Optimal Sensor Placement for Infrastructure System Monitoring using Probabilistic Graphical Models and Value of Information

Submitted in partial fulfillment of the requirements for

the degree of

Doctor of Philosophy

in

Civil and Environmental Engineering

Carl Albert Malings

B.Eng., Civil Engineering, The Cooper Union for the Advancement of Science and Art

M.S., Civil and Environmental Engineering, Carnegie Mellon University

Carnegie Mellon University

Pittsburgh, PA

May 2017

Copyright © 2017 Carl Malings

All Rights Reserved

Acknowledgements

I would like to acknowledge the support I received from many sources over the course of my doctoral studies. First, I would like to thank my advisor and doctoral committee chair Dr. Matteo Pozzi for his guidance and support during my research. By providing me with advice, guidance, financial support, and constructive criticism, he allowed me to better define my research interests and objectives and to more clearly and directly express my ideas and conclusions through scientific publications and conference presentations. He also provided me with valuable experience concerning the logistical challenges of research, including securing funding and disseminating research results. Second, I would like to thank the other members of my thesis committee (Dr. Andreas Krause, Dr. Kelly Klima, and Dr. H. Scott Matthews) for their guidance and feedback on the substance and direction of my research. Their comments during my thesis proposal and afterwards helped me to refine my research objectives and identify new approaches to addressing potential problems.

The primary funding source for this research was the Department of Civil and Environmental Engineering at Carnegie Mellon University. Additionally, the first year of research, including funding for computational resources, was provided for by the College of Engineering Dean's Fellowship. I would like to thank both the college and the department for providing this crucial initial support, as well as for the continuing financial support provided by the department during subsequent years of the research.

With regards to the design and deployment of the SHM system for Scott Hall, I would like to thank Enze Li, Xin Fu, Guangshi Shan, Chen Shou, Jing Yang, Yixiong Bao, Xuan Li, and Zitian Zhang for their assistance with the project. I would also like to acknowledge the invaluable support of Micron Optics, Inc., especially that of Michael Baez and Tiana Ruchti, throughout the fiber optic strain gauge sensor installation and data collection process, as well as that of Civionics, Inc., especially that of Andrew Zimmerman, for assistance with wireless sensor platforms. I would also like to acknowledge the outstanding assistance and patience of the Scott Hall construction team, including project manager Maxim Dorosa, Jendoco Construction Corp., Stantec Architects, Arup Steel Fabricators, and Ruthrauff-Sauer Mechanical Contractors during the sensing network design process, and for performing the installation of the sensors themselves during the construction phase. Finally, I would like to thank Irem Velibeyoglu and Matteo Pozzi for assistance with data collection and processing.

For work relating to the SHADE project, I would like to thank the other members of the project team: Matteo Pozzi, Kelly Klima, Mario Bergés, Elie Bou-Zeid, and Prathap Ramamurthy. Specifically, I would like to thank Prathap Ramamurthy for furnishing the urban temperature simulation data and answering my questions related to simulation methods and outputs, Elie Bou-Zeid for an introduction to urban meteorology and for many useful discussions on the subject, Kelly Klima for discussions related to population vulnerability and the policy issues relating to extreme temperatures, and Mario Bergés and Matteo Pozzi for general consultations

on computational approaches, sensing, and presentation of the project results. This project also received financial support from various sources. This work was funded in part by the Metro21 Initiative at Carnegie Mellon University. Elie Bou-Zeid was funded by the Helen Shipley Hunt Fund of Princeton University and by US National Science Foundation's Sustainability Research Network Cooperative Agreement # 1444758. This work was also funded in part by the Dowd Fellowship from the College of Engineering at Carnegie Mellon University. The authors would like to thank Philip and Marsha Dowd for their financial support and encouragement. The authors would also like to thank S. Harris for contribution of the project title and acronym.

I would like to thank the faculty and staff of the Civil and Environmental Engineering department at CMU for their instruction during my coursework and their assistance with technical and logistical issues. I would especially like to thank Professors VanBriesen, Thompson, and Samaras for their guidance during my tenures as Teaching Assistant for the Introduction to Civil and Environmental Engineering course. I would like to thank the department staff for their invaluable assistance with logistical issues. I would also like to thank my classmates and colleagues at CMU, especially my office-mates, for their advice, feedback, and assistance, as well as for stimulating on- and off-topic discussions.

I would like to thank my family and my significant other for their patience and support during this process. Without their understanding, encouragement, and backing, this work would not have been possible.

Finally, I would like to thank the reader for taking an interest in this body of work. I hope this document will provide a clear, comprehensive, and thorough description of issues and approaches related to the optimization of information gathering to support the management of infrastructure systems.

Abstract

Civil infrastructure systems form the backbone of modern civilization, providing the basic services that allow society to function. Effective management of these systems requires decision-making about the allocation of limited resources to maintain and repair infrastructure components and to replace failed or obsolete components. Making informed decisions requires an understanding of the state of the system; such an understanding can be achieved through a computational or conceptual system model combined with information gathered on the system via inspections or sensors. Gathering of this information, referred to generally as sensing, should be optimized to best support the decision-making and system management processes, in order to reduce long-term operational costs and improve infrastructure performance.

In this work, an approach to optimal sensing in infrastructure systems is developed by combining probabilistic graphical models of infrastructure system behavior with the value of information (VoI) metric, which quantifies the utility of information gathering efforts (referred to generally as sensor placements) in supporting decision-making in uncertain systems. Computational methods are presented for the efficient evaluation and optimization of the VoI metric based on the probabilistic model structure. Various case studies on the application of this approach to managing infrastructure systems are presented, illustrating the flexibility of the basic method as well as various special cases for its practical implementation.

Three main contributions are presented in this work. First, while the computational complexity of the VoI metric generally grows exponentially with the number of components, growth can be greatly reduced in systems with certain topologies (designated as cumulative topologies). Following from this, an efficient approach to VoI computation based on a cumulative topology and Gaussian random field model is developed and presented. Second, in systems with non-cumulative topologies, approximate techniques may be used to evaluate the VoI metric. This work presents extensive investigations of such systems and draws some general conclusions about the behavior of this metric. Third, this work presents several complete application cases for probabilistic modeling techniques and the VoI metric in supporting infrastructure system management. Case studies are presented in structural health monitoring, seismic risk mitigation, and extreme temperature response in urban areas. Other minor contributions included in this work are theoretical and empirical comparisons of the VoI with other sensor placement metrics and an extension of the developed sensor placement method to systems that evolve in time.

Overall, this work illustrates how probabilistic graphical models and the VoI metric can allow for efficient sensor placement optimization to support infrastructure system management. Areas of future work to expand on the results presented here include the development of approximate, heuristic methods to support efficient sensor placement in non-cumulative system topologies, as well as further validation of the efficient sensing optimization approaches used in this work.

Executive Summary

The management of civil engineering infrastructures involves allocating limited resources to where they are most needed to improve or maintain the functionality of the system in question. Decisions about this allocation are made using information about the consequences of failures in different components of the system, awareness of the possible actions that can be taken to avoid these failures or mitigate their consequences, and knowledge about the current state of the system, i.e., the chances of different system components failing in the near future. Knowledge about the system state comes from a combination of expert assessments and past experience about the behavior of the system or of similar systems, and about the various situations that might cause component and system failures. This knowledge is also updated by gathering information within the system, e.g., by performing inspections on components to determine their state or gathering data on environmental factors that determine the pressures placed on the system. This additional collected information helps guide decision-making for management of the system, with the overall goal being to reduce the cost of managing the system by selecting appropriate maintenance activities to avert potential failures, while avoiding unnecessary activities.

In current practice, the collection of information about a system (when it is performed) is often done according to prescribed plans that do not necessarily acquire the most timely and useful information about the system, nor use this information in the most effective way. For some systems, information is not collected at all, and maintenance is performed according to fixed prior plans. Such plans can lead to an inefficient allocation of resources, as some components are maintained or replaced while still being in good condition, whereas others that degrade more rapidly than anticipated may not be replaced in time to prevent a costly failure. When information is collected, this is typically done according to a fixed plan where all components in the system are inspected at a certain interval. This places a high demand on the inspectors of the system, and in many cases shortages of qualified inspectors lead to these plans not being followed, or to inspections being made in a cursory manner. Prioritization of these inspections, based on the presumed usefulness of the inspection results in supporting better decision-making, is a way in which this burden can be reduced while still acquiring useful and timely information. Furthermore, the information collected by inspections is typically used only to update the status of the individual component being inspected. In general, this information can be combined with knowledge about how components interact and the similarities between them to provide insight into the status of the system as a whole.

The main contributions of this work are to develop and demonstrate methods for prioritizing the collection of information to support the management of infrastructure, to alleviate the problems discussed above. The proposed methods involve two key elements. First, probabilistic models of the system allow for quantification of the uncertainties associated with system management, and also for the modeling of similarities within the system, such that information collected in a

specific area can be used to update knowledge about the behavior of the system as a whole. The general approach to defining these models is discussed in Chapter 2. Second, the value of information (VoI) metric is used to prioritize the inspection of the system. The VoI is a measure of how much additional information is worth to a decision-maker; it quantifies how much the management cost of the system is expected to be reduced by having additional information available to support better decision-making. This allows for a quantitative, monetary value to be assigned to different sources of information, and therefore for the identification of information that has the biggest potential impact in terms of making better decisions to reduce system management costs. The cost-effectiveness of alternative sources of information can also be assessed by the VoI. This metric is introduced in Chapter 4, where an in-depth analysis of its properties and a comparison with other possible metrics for information assessment are provided.

Unfortunately, VoI can be difficult to evaluate in large systems. This is due to many factors, including potentially complicated interactions between the states of various components in how they operate as part of the system, as well as interactions between decisions made to manage different parts of the system that have synergistic effects on the system as a whole. In this work, a key contribution is to identify different types of systems for which the VoI can be more easily evaluated. Chapter 3 discusses various types of systems considered, and Chapter 5 provides details on how the computation of VoI can be made simpler in these cases. Furthermore, Chapter 7 presents special considerations for evaluating VoI in systems that evolve over time, where choices must be made about both where and when to gather information.

The primary advantage of using VoI to prioritize the collection of information is its ability to quantify the improvement in decision-making outcomes due to the availability of additional information. This allows for informed trade-offs to be made between the costs of acquiring information and the benefits this information is likely to bring. This impact will be felt most keenly in scenarios where the cost of collecting information is high and the resources for doing so are tightly constrained, e.g., where there is a shortage of qualified human inspectors or where expensive measurement equipment is required. The contributions presented in this work will help system managers realize this benefit by being able to more readily calculate VoI for a variety of civil infrastructures.

A few shortcomings of VoI-based sensing optimization must also be acknowledged. First, in situations where information gathering is inexpensive, e.g., where many low-cost sensors can be easily deployed to monitor the system, the additional effort needed to optimally allocate these measurements throughout the system may not be justified. Furthermore, in situations where the underlying performance of the system (including the consequences of failure of components and the possible actions that can avoid these failures) are poorly understood, the VoI cannot be evaluated in a meaningful way. Finally, in certain situations, efficient methods presented in this work to prioritize information collection based on VoI can have poor performance. These situations, as well as potential workarounds for some of these shortcomings, are presented briefly in Chapter 11.

The final major contribution of this work is to present a variety of different application case-studies to relevant problems in civil and environmental engineering. First, in Chapter 6, examples are presented in simplified one-dimensional (representing a beam subjected to loading) and two-dimensional (representing a metal surface subjected to corrosion) systems. These examples illustrate how the VoI metric conforms well to several intuitive ideas of the relevance of information; specifically, the VoI is high when the information is precise and accurate, when it is unclear what action(s) will be best to manage the system, and/or when the consequences of making an incorrect decision are relatively severe. Three more in-depth case studies in realistic system are also provided. In Chapter 8, the monitoring of loading in a building is considered. An optimized subset of strain gauges installed on the new Scott Hall building at Carnegie Mellon University is identified that best supports management of the structure under potential differential settlement of its foundations. Chapter 9 presents an application to seismic risk assessment in the San Francisco Bay area. This application involves the collection of information after an earthquake in order to make appropriate decisions about the closure of bridges and tunnels for safety reasons. An example on the prediction of and response to summer temperatures in the city of Pittsburgh is presented in Chapter 10. Here, information from various sources, including regional weather forecasts, climate simulations, and local temperature gauges, is combined to support decision-making about the issuance of heat advisories to avoid the consequences of extreme heat exposure. Together, these examples illustrate the adaptability and wide applicability of the methods and metrics for sensor placement developed in this work.

In summary, this work develops and demonstrates methods for the optimal collection of information to support the management of infrastructures. The VoI metric is used to support this, and details are provided about how this metric can be computed efficiently in large systems under certain assumptions. By directly quantifying the benefits of information in monetary terms, this metric allows for better allocation of limited inspection resources throughout systems, as opposed to current rigid prescriptive plans for information gathering. In illustrating the application of these methods to various problems of interest in civil and environmental engineering systems management, it is hoped that the benefits of this approach will be made clear, and that appropriate, optimized decisions about information gathering can be made to support the management of other infrastructures that are critical to modern society.

Table of Contents

Chapter 1 Introduction **1**

1.1. Motivation 1
1.2. Background 4
1.2.1. Current Infrastructure Monitoring Practices *4*
1.2.2. Structural Health Monitoring *5*
1.2.3. Infrastructure System Modeling *6*
1.2.4. Optimal Sensor Placement *10*
1.2.5. Value of Information *11*
1.2.6. Summary of Research Gaps *14*
1.3. Research Questions 15
1.4. Key Results and Contributions 16
1.5. Overview of the Document 17

Chapter 2 Problem Statement **20**

2.1. Basic Definitions 22
2.1.1. Random Fields *22*
2.1.2. Observations *24*
2.1.3. States *25*
2.1.4. Actions *27*
2.1.5. Losses *27*
2.2. Optimal Sensor Placement 28
2.2.1. Optimization Approach *29*

Chapter 3 Modeling Infrastructure Systems **33**

3.1. Probabilistic Graphical Modelling Approach 33
3.2. Infrastructure System Topology 35
3.2.1. Independent Systems *35*
3.2.2. Cumulative Systems *37*
3.2.3. Non-Cumulative Systems *38*
3.3. Gaussian Random Field Infrastructure Model 39
3.3.1. Gaussian Random Fields *40*
3.3.2. Linear Observations *40*
3.3.3. Affine Limit States *42*
3.3.4. Binary Component States *43*
3.3.5. Binary System States *44*

Chapter 4 The Value of Information Metric **45**

4.1. Value of Information 45

4.1.1. Bounds on the Value of Information *47*
4.1.2. Regret-based Formulation *47*
4.1.3. Sensor Placement using Value of Information *48*
4.2. Prototypical Loss Functions 49
4.2.1. Cumulative Systems *50*
4.2.2. Non-Cumulative Systems *50*
4.3. Parametric Analysis on a Single-Component System 51
4.4. Relationship to Other Sensor Placement Metrics 57
4.4.1. Prediction Error *58*
4.4.2. Misclassification Rate *60*
4.4.3. Conditional Entropy *61*
4.5. Non-Submodularity of the Value of Information Metric 66

Chapter 5 Metric Evaluation and System Topology 69

5.1. Computational Difficulties 69
5.2. Decomposable Loss Functions 72
5.2.1. Efficient Evaluation for Gaussian Prediction Error *73*
5.2.2. Efficient Evaluation for Binary Components *75*
5.3. Non-Decomposable Loss Functions 77
5.3.1. General Metric Evaluation via Sampling *77*
5.3.2. Approximate Techniques for Series and Parallel Systems *78*
5.4. Investigations in Five Component Systems 80

Chapter 6 Comparative Examples for Sensor Placement 86

6.1. One-Dimensional Example in Structural Loading 86
6.1.1. Cumulative System *89*
6.1.2. Global System *90*
6.1.3. Series and Parallel Systems *91*
6.1.4. Multiple Limit-State Functions *93*
6.2. Two-Dimensional Example in Corrosion Prevention 96
6.2.1. Cumulative System *98*
6.2.2. Series System *100*
6.2.3. Global Action *101*
6.2.4. Comparison to Random Field Estimation *102*
6.2.5. Comparison to Online Sensor Placement *102*

Chapter 7 Extension to Spatio-Temporal Systems 105

7.1. Formulation of Spatio-Temporal Value of Information 106
7.1.1. Temporally Decomposable Loss Functions *109*
7.1.2. Offline and Online Sensor Placement and Scheduling *112*
7.2. Demonstrative Examples 114

7.2.1. Classification in an Evolving 1-D Random Field *114*
7.2.2. Management of a River Water Intake System *115*
7.2.3. Differential Settlement in a Foundation System *117*

Chapter 8 Structural Health Monitoring of Scott Hall 122

8.1. Overview of Scott Hall and its Instrumentation 122
8.2. Probabilistic Strain Field Modeling 127
8.3. Sensor Placement Results 130
8.4. Sensor Placement and Scheduling Results 131

Chapter 9 Seismic Risk in the San Francisco Bay Area 133

9.1. Seismic Risk Model 134
9.1.1. Value of Information in Poisson Processes *136*
9.2. Application to San Francisco Bay Area 138
9.2.1. Value of Information Density *142*

Chapter 10 Urban Heat Risk Prediction and Response 145

10.1. Background on Urban Heat Risk and Response 146
10.2. Probabilistic Temperature Modeling 149
10.2.1. Definition *149*
10.2.2. Measurements *153*
10.2.3. Calibration *155*
10.2.4. Validation *164*
10.3. Temperature Monitoring 172
10.3.1. Risk Formulation *173*
10.3.2. Sensing Metrics Considered *176*
10.4. Sensor Placement Results 179
10.4.1. With Global Mean and Cyclic Pattern Measures *179*
10.4.2. Without Global Mean and Cyclic Pattern Measures *187*

Chapter 11 Open Issues and Future Work 194

11.1. Performance of the Greedy Optimization Approach 194
11.2. Efficient Metric Evaluation in Non-Cumulative Systems 202
11.3. Uncertainty in Probabilistic Model Definitions 204
11.4. Additional Application Cases 205

Chapter 12 Conclusions 208

12.1. Related Publications 213

References 215

Appendix A: Value of Information Bounds and Relationships 228

A.1. Lower Bound for the Value of Information 228

A.2. Upper Bound for the Value of Information 229
A.3. Regret-based and Loss-based Formulation Equivalence 231
A.4. Greedy Online versus Offline Sensor Placement..... 231

Appendix B: Submodularity Investigation..... 234

B.1. Conditional Entropy 234
B.2. Mutual Information 237
B.3. Value of Information..... 238

Appendix C: Parameters for San Francisco Bay Area Example 242

List of Tables

Table 0-A: Font styles and their common meanings. xix
Table 0-B: Latin alphabet characters and their common meanings. xx
Table 0-C: Greek alphabet characters and their common meanings. xxi
Table 0-D: Special characters and their common meanings. xxii
Table 0-E: Common abbreviations. xxii
Table 6-A: VoI metric results for the systems of Sections 6.1.1 through 6.1.3. 90
Table 9-A: Sensor placements by the conditional entropy metric, in order of greedy selection. 140
Table 9-B: Sensor placements by the VoI metric, in order of greedy selection. 141
Table 10-A: WRF-PUCM simulation settings. 158
Table 10-B: Calibrated parameters of the Gaussian process model for the global mean temperature. 159
Table 10-C: Calibrated parameters of the covariance function for the cyclic temperature pattern. 161
Table 10-D: Calibrated parameters of the covariance function for the residual temperature field. 162
Table 10-E: Test cases for model validation. 165
Table 10-F: Comparative results for New York City, listing RMSE for different model cases and prediction lead times. 168
Table 10-G: Comparative results in Pittsburgh, listing RMSE for different model cases and prediction lead times. 171
Table 10-H: Comparative results for sensor placement metrics. 185
Table 10-I: Comparative results for VoI of various measurement sets. 191
Table 11-A: Randomized system generation method settings. 195
Table C-A: List of infrastructure components. 242
Table C-B: List of component capacity parameters. Based on models from HAZUS (2013). .. 243
Table C-C: List of component capacity correlations. 244

List of Figures

Figure 3-1: Generic influence diagram of an infrastructure system and its management. 34
Figure 3-2: Simplified influence diagram, modeling a loss function of the form L(**f**,**a**). 34
Figure 3-3: Influence diagram for a system with independent components monitored and managed independently. 36
Figure 3-4: Influence diagram for a system with a cumulative topology. 37
Figure 3-5: Influence diagram for a system with a non-cumulative topology. 38
Figure 4-1: a) Example expected loss versus failure probability for management of a single component; b) Example expected regret versus failure probability for management of a single component; c) Example prior and posterior failure probability distributions for a single component. 53
Figure 4-2: Parametric analysis results for VoI in a single-component system, with cross-sections for a) $r = 0$; and b) $\beta = 2$. 56
Figure 4-3: Misclassification rate as a function of failure probability for a binary component. .. 61
Figure 4-4: Entropy as a function of failure probability for a binary component. 63
Figure 5-1: a) Prior expected loss as a function of correlation for various system topologies; b) VoI scaled by the prior expected loss as a function of number of measurements for various system topologies and correlation coefficients. 82
Figure 5-2: a) Computational time for VoI evaluation versus number of components for several system topologies, with dotted line indicating exponential growth rate; b) Extended plot of computation time versus number of components for cumulative systems, with dotted line indicating cubic growth rate. 85
Figure 6-1: Prior one-dimensional Gaussian process loading model, with an example instantiation of the loading field. 87
Figure 6-2: Posterior one-dimensional Gaussian process loading model, conditional to an example observation. 88
Figure 6-3: VoI metric results for a cumulative system topology. 89
Figure 6-4: VoI metric results for a global system topology. 91
Figure 6-5: VoI metric results for a) series and b) parallel system topologies. 92
Figure 6-6: Influence lines for structural loading. 93
Figure 6-7: VoI metric results considering failure conditions separately, for a) estimating limit-state variables separately and b) managing the system under each failure condition. 94
Figure 6-8: VoI metric results considering failure conditions together, for a) estimating the limit-state variables together and b) managing the series system of failure conditions. 95
Figure 6-9: a) Prior mean of the random demand field; b) Prior standard deviation of the random demand field; c) Prior probability of corrosion for each location. 96
Figure 6-10: a) Realization of the demand field; b) Realization of the capacity field; c) Realization of the corrosion field (i.e., the component states), with yellow denoting corrosion. 97

Figure 6-11: a) Prior expected loss for the management of the ceiling as a cumulative system; b) Prior optimal actions for the management of the ceiling, with yellow corresponding to a choice to apply coatings locally to prevent corrosion. 99
Figure 6-12: Placements for five humidity sensors optimized offline by the VoI metric for a cumulative system (left) and the VoI provided by these sensors (right). 99
Figure 6-13: Placements for five humidity sensors optimized offline by the VoI metric in a series system of the lower-right quadrant (left) and the net VoI provided by these sensors using three alternative cost functions (right). 100
Figure 6-14: Placements for five humidity sensors optimized offline by the VoI metric considering a global management action (left) and the VoI provided by these sensors (right).. 101
Figure 6-15: Placements for five humidity sensors optimized offline by the random field prediction metric (left) and the VoI and posterior random field variance reduction provided by these sensors (right). 102
Figure 6-16: Placements for five humidity sensors optimized online by the VoI metric for a cumulative system (left) and the reduction in loss due to improved posterior actions (right).... 103
Figure 7-1: Optimal placement and scheduling of ten measurements in a one-dimensional system for classification. 115
Figure 7-2: Optimal placement and scheduling of five water contamination samples along a 10km stretch of river to support water intake system management, plotted above the contamination level uncertainty reduction. 117
Figure 7-3: a) Layout of columns subjected to settlement; b) Optimal offline sensing scheme for this system, with measures denoted by x's; c) VoI for the scheme, with respect to the number of sensors greedily selected, with the optimal set indicated by the peak of the net VoI; d) Example settlement of column 1, with the prior 95% confidence interval (gray area), actual settlement (grey line), and settlement predicted by the optimal measurement scheme (dotted line)..... 119
Figure 7-4: a) Comparison of forward greedy optimized (o) and reverse greedy optimized (x) sensor placement and scheduling schemes in the case of unknown measurement offsets; b) Comparison of the net VoI provided by sets of various seized optimized via forward and reverse greedy approaches; the vertical dotted line indicates the schemes depicted in (a). 121
Figure 8-1: Computer rendering of Scott Hall's North Wing structure, between the existing Wean Hall (left) and Hammerschlag Hall (right). Image from http://www.cmu.edu/cdfd/scott-hall/.. 123
Figure 8-2: Micron Optics fiber-optic strain gauge (a) and interrogation unit (b). Images from www.micronoptics.com. 124
Figure 8-3: Diagram of strain gauge installation on structural columns. 125
Figure 8-4: Photographs of installed sensors on column TC-E from the northwest (a) and southwest (b). 126
Figure 8-5: Diagram of sensors, columns, and fiber optic cable configuration. 127
Figure 8-6: Example of sensor data updating the prior distribution over relative displacement trends for a time interval to a posterior distribution. The true trend for the time interval computed from the data is indicated for reference. 129

Figure 8-7: Optimized sensor placements (left) and associated VoI (right). 130
Figure 8-8: Results for sensor placement (a), scheduling (b), and placement and scheduling (c). 131
Figure 9-1: Proposed sensor placements based on the conditional entropy metric. Background image from www.maps.google.com. 139
Figure 9-2: Proposed sensor placements based on the VoI metric. Background image from www.maps.google.com. 140
Figure 9-3: Net VoI versus number of sensors for sensing schemes optimized via the conditional entropy and VoI metrics, and for an example of random sensor placement. 141
Figure 9-4: Density of the discounted cumulative VoI over earthquake magnitude for the proposed scheme of Figure 9-2 (blue), and for all possible sensors (black). The probability density function for earthquake magnitudes is also presented for reference. 143
Figure 9-5: Density of the discounted cumulative VoI over earthquake epicenter locations for the proposed scheme of Figure 9-2. Note that the color scale is logarithmic. Background image from www.maps.google.com. 144
Figure 10-1: Satellite imagery of the simulation regions for New York City (a) and Pittsburgh (b). Satellite images from www.maps.google.com. 158
Figure 10-2: Comparison between the empirically observed temporal correlation of the global mean temperature, as computed from simulated temperature data for New York City (dotted line), and the calibrated covariance function of Equation 10-5 (solid line). 159
Figure 10-3: Temperature predictions using a Gaussian process model for the global mean temperature. The average temperatures before 6/22 are used as training data, and the resulting 95% confidence bounds on the predictions of future temperatures are shown as the grey area, along with the actual temperature record (black line). 160
Figure 10-4: Comparison between the empirically observed spatial (a) and temporal (b) correlations of the cyclic temperature pattern, as computed from simulated temperatures for New York City (dotted line), and the calibrated covariance functions of Equations 10-8 and 10-9 (solid line). 160
Figure 10-5: Posterior mean of the cyclic temperature patterns for New York City at local midnight (a) and noon (b), conditioned on measurements of these patterns, as in Equation 10-18, obtained from the training data set. 161
Figure 10-6: Comparisons between the empirical spatial (a) and temporal (b) correlations in the residual temperature fields for New York City, NY (dotted line), compared to the trained covariance functions of Equations 10-13 and 10-14 (solid line). 162
Figure 10-7: Spatial standard deviation in the residual temperature field for the region of New York City, NY (a) and Pittsburgh, PA (b). 163
Figure 10-8: Comparison of the temperature record from WRF (black line) for a location on Manhattan Island with the 95% confidence interval (grey area) of the posterior model for one-hour-ahead prediction under use case 4. 166

Figure 10-9: Predictive performance of the probabilistic model for New York City, NY, under various use cases. For comparison, the performances of two baseline prediction cases are also provided. 167
Figure 10-10: Quantile-quantile plot of 1,000 empirical residuals of the probabilistic model for New York City, NY, under each use case. The 95% confidence interval of the Gaussian distribution is indicated between the vertical dashed lines. The grey dashed line indicates the ideal normal distribution. 169
Figure 10-11: Results for posterior model performance under use case 4 with different levels of data (as measured by % Samples, or the fraction of discrete locations at which data are collected for updating the model) for New York City, NY. 170
Figure 10-12: Comparison of actual temperature in downtown Pittsburgh from the simulated temperature data of the validation period on the date indicated (solid line) with the measurement of the nearest weather station (dotted line) and the 95% confidence interval of the posterior prediction of the Gaussian process temperature model (grey area) for one-hour-ahead prediction under case 4 with local data from the single weather station (a) and using these local measurement data as well as prior measures of the cyclic pattern obtained from training data (b). 172
Figure 10-13: Vulnerability index map for Pittsburgh, PA. 174
Figure 10-14: An example risk index based on 95th percentile peak daily temperature in Pittsburgh, PA, and the vulnerability map of Figure 10-13. This index is normalized between 0 and 1. 176
Figure 10-15: Optimized measurement locations for Pittsburgh, PA, supporting the reduction of conditional entropy of the temperature field. Sensing locations are indicated by x's, while the background indicates the posterior standard deviation of the temperature field conditional to these measurements. Numbers next to the sensing locations indicate the order the placements were selected by the greedy optimization algorithm. The city's outline is also indicated. 180
Figure 10-16: Reduction in entropy resulting from the sensor placements through greedy optimization (solid line). The line with slope corresponding to 82% of the reduction of the first sensor is also plotted (dashed line), and the point at which this line is tangent to the entropy reduction curve identifies the optimal number of sensors (in this case 9). 181
Figure 10-17: Optimized measurement locations for Pittsburgh, PA, supporting the reduction of prediction error of the temperature field weighted by the vulnerability field. Sensing locations are indicated by x's, while the background indicates the posterior weighted prediction error conditional to these measurements. Numbers next to the sensing locations indicate the order the placements were selected by the greedy optimization algorithm. The city's outline is also indicated. 182
Figure 10-18: Gain in weighted prediction error resulting from the sensor placements through greedy optimization (solid line). The line with slope corresponding to 6% of the gain of the first sensor is also plotted (dashed line), and the point at which this line is tangent to the weighted prediction error gain curve identifies the optimal number of sensors (in this case 9). 183

Figure 10-19: Optimized measurement locations for Pittsburgh, PA, supporting decision-making for heat advisory issuance via the VoI metric. Sensing locations are indicated by x's, while the background indicates the expected regret corresponding with the indicated sensor placements, distributed over space, for June 27, 2012 at 1630 hours. Numbers next to the sensing locations indicate the order the placements were selected by the greedy optimization algorithm. The city's outline is also indicated. 184
Figure 10-20: VoI resulting from the sensor placements through greedy optimization. Note that the net VoI is maximized when $Y = 9$ for a cost of 0.025 per sensor; this is the optimal number of sensors indicated in Figure 10-19. The net VoI under other costs are also shown. 185
Figure 10-21: Comparison between optimized sensor locations under the conditional entropy metric, with (A) and without (B) global mean and cyclic pattern data available. 188
Figure 10-22: Comparison between optimized sensor locations under the weighted prediction error metric, with (A) and without (B) global mean and cyclic pattern data available. 189
Figure 10-23: Comparison between optimized sensor locations under the VoI metric, with (A) and without (B) global mean and cyclic pattern data available. 190
Figure 10-24: VoI versus prediction lead time for two measurement sets. 192
Figure 11-1: Comparative plot of the VoI provided by greedily and exactly optimized measurement sets of the same size across 100 simulated systems. The identity line (dashed) is plotted for reference. 196
Figure 11-2: Empirical cumulative distribution function for the ratio of the VoI of optimal measurement sets of the same size obtained by greedy and exact optimization. Note the scale of the horizontal axis and discontinuity at 1. 197
Figure 11-3: Comparative plot of the VoI provided by forward and reverse greedily optimized measurement sets of the same size across 100 simulated systems. The identity line (dashed) is plotted for reference. 198
Figure 11-4: Order in which measures are to be included via forward, reverse, and proposed heuristic-based greedy optimization approaches (a) and the VoI of the resulting sets (b). 201
Figure 11-5: Comparative plot of the VoI provided by sets of measurements in series systems and in systems with equivalent cumulative topologies. Results are divided into three groups: sets where the series and equivalent cumulative systems yield similar VoI values (blue), sets where negligible value is provided by sensors in the series system (black), and sets where the VoI obtained in both systems follow an unknown relationship (red). The identity line (dashed) is plotted for reference. 203

Symbols and Notation

This section describes symbolic and notational conventions adopted in this paper. These will typically be defined following their first usage in the text.

Table 0-A summarizes the common usages associated with different styles of notation (i.e., upper- or lower-case, bold, and italic styles). Generally, italics are used to denote scalar quantities and variables, non-italics to denote functions, and bold characters to denote vectors (if lower-case) or matrices (if upper-case). There may in some cases be exceptions to the typical cases presented here (for example, the use of an upper-case L to avoid ambiguity between the lower-case l and the numeral 1).

Table 0-A: Font styles and their common meanings.

	Style	**Typical Usage**	**Example**
a	lower-case, non-bold, non-italic	Function of scalar or vector	$\mathrm{f}(a)$
a	lower-case, non-bold, italic	Scalar quantity	a
$\mathbf{a}$	lower-case, bold, non-italic	Vector quantity	$\mathbf{a} = \{a_1, \dots, a_n\}$
A	upper-case, non-bold, non-italic	Function of variable or set	$\mathrm{Vol}(Y)$
A	upper-case, non-bold, italic	Name of a variable or set	X
$\mathbf{A}$	upper-case, bold, non-italic	Matrix	$\mathbf{\Sigma}_Y$

In the following paragraphs, several notational conventions used in this paper will be presented and discussed. First, to describe the probability of an event occurring, the notation $\mathbb{P}(\cdot)$ is adopted. For example, the notation $\mathbb{P}(X = x)$ indicates the probability that a random variable X will take on scalar value x. To determine this, the probability distribution for the variable X is denoted as p_X; the notation $X \sim \mathrm{p}_X$ indicates that this variable is distributed according to this distribution. Where such a distribution has been defined, the probability (or probability density, for continuous variables) of variable X taking on value x is equivalently denoted as $\mathrm{p}_X(x)$, i.e., the evaluation of the function p_X at value x. The same notational convention holds for vector-valued variables, i.e., $\mathbb{P}(X = \mathbf{x}) = \mathrm{p}_X(\mathbf{x})$.

Following these definitions, a statistical expectation is denoted as follows:

$$\mathbb{E}_X \mathrm{f}(x) = \int_{x \in \mathrm{Dom}(X)} \mathrm{f}(x)\, \mathrm{p}_X(x)\, \mathrm{d}x \tag{0-1}$$

That is, the statistical expectation of function $\mathrm{f}(x)$ with respect to variable X is denoted $\mathbb{E}_X \mathrm{f}(x)$, and is computed as the integral, over all values x in $\mathrm{Dom}(X)$, the domain of variable X, of the product of function $\mathrm{f}(x)$ evaluated at value x and the probability of this value under the distribution for variable X, $\mathrm{p}_X(x)\, \mathrm{d}x$.

Finally, an optimization problem is typically denoted as follows:

$$x^* = \text{argmax}_{x \in \text{Dom}(X)} \, \text{f}(x) \tag{0-2}$$

That is, the value of variable X that maximizes function $\text{f}(x)$ is x^*. The same formulation is used for set functions:

$$A^* = \text{argmax}_{A \in \mathcal{A}} \, \text{F}(A) \tag{0-3}$$

That is, the set A that maximizes the value of set function $\text{F}(A)$ is A^*. The notation $\mathcal{A}$ refers to the powerset (i.e., the collection of feasible sets) from which the set A is to be selected.

The following pair of tables list capital and lower-case letters in the Latin and Greek alphabets, and the definitions that are typically associated with these variables in the paper.

Table 0-B: Latin alphabet characters and their common meanings.

a	action	A	set of actions
b	budget	B	
c	capacity	C	cost
d	demand	D	distance
e	natural logarithm base	E	expectation (prefix)
f	random field (scalar or vector of values)	F	random field (set)
g	Limit-state	G	set of limit-states
h	indexing value	H	entropy
i	indexing value	I	identity
j	indexing value	J	
k	covariance function	K	
l		L	loss
m	mapping	M	sensor placement objective
n	number (without subscript: number of components)	N	
o		O	order notation
p	probability (distribution)	P	probability (value)
q		Q	
r	measurement noise ratio	R	regret or risk
s	state	S	set of states
t	time	T	temperature
u		U	
v	velocity	V	vulnerability
w	model parameter	W	model parameters (set)
x	spatial coordinate	X	set of spatial coordinates
y	observation	Y	set of observations
z		Z	weight matrix

Table 0-C: Greek alphabet characters and their common meanings.

α		Α	see A
β	reliability index	Β	see B
γ	discount factor	Γ	
δ	Kronecker delta function	Δ	change (prefix)
ϵ	error	Ε	see E
ζ		Ζ	see Z
η		Η	see H
θ		Θ	
ι		Ι	see I
κ		Κ	see K
λ	correlation length	Λ	recurrence rate
μ	mean or mean function	Μ	see M
ν	unused due to similarity with v	Ν	see N
ξ		Ξ	limit-state function matrix
o	see o	Ο	see O
π	ratio of circle circumference to diameter	Π	
ρ	correlation	Ρ	see P
σ	standard deviation	Σ	covariance matrix
τ	set of time steps	Τ	see T
υ	unused due to similarity with v	Υ	unused due to similarity with Y
ϕ		Φ	Gaussian cumulative distribution function
χ		Χ	see X
ψ		Ψ	
ω		Ω	observation matrix

Table 0-D lists several special characters that are used along with their typical definitions.

Table 0-D: Special characters and their common meanings.

$\mathbb{E}$	statistical expectation (with respect to subscripted variable)
$\mathbb{I}$	indicator function (value 1 when argument is true, otherwise 0)
$\mathbb{P}$	probability of an event
$\mathbb{S}$	system state (function)
$\mathbb{s}$	system state (variable)
$\mathcal{A}$	powerset of possible actions
$\mathcal{GP}$	Gaussian process
$\mathcal{N}$	Gaussian distribution
$\mathcal{T}$	temporal domain
$\mathcal{X}$	spatial domain
$\mathcal{Y}$	powerset of potential observations
$\emptyset$	empty set
$\cup$	set union
$\setminus$	set subtraction
$\in$	element of a set

Finally, Table 0-E lists common abbreviations used in this work.

Table 0-E: Common abbreviations.

ASCE	American Society of Civil Engineers
NWS	National Weather Service
PGA	peak ground acceleration
PGM	probabilistic graphical model
RH	relative humidity
SHM	structural health monitoring
USGS	United States Geological Survey
VoI	value of information
VoPI	value of perfect information

Chapter 1

Introduction

This chapter introduces the work, providing the motivation for optimal sensor placement to support infrastructure system monitoring and management. Background information on prior work in this area is summarized, providing context for the contributions presented here. Based on this background and identified research gaps, one general research question and several specific sub-questions are proposed for this work, and significant contributions related to these questions are briefly summarized. Finally, the organization of the remainder of the document is discussed.

1.1. Motivation

Civil engineering infrastructure systems, such as water supply, wastewater collection, transportation, and power distribution systems, are essential to the smooth functioning of modern society. These infrastructure systems are made up of numerous components (e.g., transportation systems may include bridges, roadways, and/or tunnels) that are interconnected with each other and together constitute the system. Infrastructure systems are also interdependent in that the proper functioning of components within one system can depend on the operation of another system (e.g., pumping stations in a water distribution system depend on electricity from a power distribution system). Components of infrastructure systems are subject to both gradual degradation over time as well as extreme hazard events (e.g., earthquakes or hurricanes). Gradual degradation, hazard events, or a combination of both can cause the failure of individual components within these systems and therefore, through the interactions of components within and between systems, failures of the infrastructure system as a whole to fulfill its function. Such failures result in significant economic and societal impacts, and potentially loss of life (ASCE, 2013, 2016). To avoid such costly infrastructure system failures, it is necessary for the managers of these systems to develop a comprehensive understanding of their behavior, taking their many interactions and interdependencies into account, and, making use of this understanding, to decide upon appropriate management actions that will reduce the potential risk or impact of system failures, subject to constraints on resources.

This task of understanding and effectively operating modern interdependent infrastructure systems is complicated by many factors. The interconnections and relationships between components in an infrastructure system are often difficult to define and model in a quantitative

way (e.g., Rinaldi et al., 2001; Buldyrev et al., 2010). The loadings and hazards to which infrastructure components will be subjected over their lifetimes are uncertain, and this uncertainty is only exacerbated by the predicted non-stationarity of future climate, i.e., climate change (e.g., Mokrech et al., 2012; Garré and Friis-Hansen, 2013; Špačková and Straub, 2015). In addition, the response of components of infrastructures to loadings and hazards are also uncertain, due to inherently imperfect knowledge of the as-built conditions of infrastructures and of their precise material and mechanical properties. Infrastructure management decision-makers must take these many sources of uncertainty into account when developing optimal management strategies for infrastructure systems.

Much research has therefore been devoted both to developing appropriate models of these uncertainties (e.g., Faber et al., 2011) and to designing non-destructive sensing and monitoring technologies to reduce these uncertainties (e.g., Büyüköztürk et al., 2013). Inspections of infrastructure components and, more recently, the development and deployment of structural health monitoring (SHM) systems to continuously provide information on the status of infrastructure components both seek to reduce the uncertainty in component properties and responses, as well as identify true loading conditions (Farrar, 2013). At the same time, modeling techniques such as the use of probabilistic graphical models (PGMs) and network-based models of interconnected infrastructure systems aim at capturing the relationships between components of infrastructure systems, allowing for quantification of both the uncertainties and the interdependencies in the behaviors of the system (Koller and Friedman, 2009). Advances in computing power have allowed the practical application of these models to infrastructure systems of a realistic scope and scale (e.g., Baker et al., 2015; González et al., 2015).

An open area for future research, and one that motivates this work, is the development of an approach that unifies the modeling of civil infrastructure performance with optimal inspection, monitoring, and decision-making for managing these systems. Such an approach will allow for the simultaneous optimization of infrastructure system monitoring and management, using metrics that quantify the utility of these activities to the managing agents of the infrastructure and to the general public. This approach will have applications in the design of inspection schemes and/or the deployment of SHM networks over infrastructure systems, allowing sensing resources to be allocated where they will have the most benefit in terms of providing useful information to improve management decision-making to reduce life-cycle system costs. This research seeks to develop such an approach, as well as demonstrate its application to a number of realistic infrastructure system management problems.

The following motivating example raises questions about sensor placement that this work will attempt to address. Consider that a structure is being subjected to an unknown load that has the potential to cause failures to different parts of the structure, with associated consequences. Reinforcements can be applied to the structure to guard against failure, at a certain cost. A prior model for the unknown loading has been developed that describes what the loads are expected to be, and also quantifies the uncertainty in these estimates via a probabilistic model. Sensors are

available that can accurately measure the load, but due to their high price, only a single sensor can be deployed. The question is: where should this sensor be placed?

An intuitive response is that the load should be measured where the potential loading is the most variable, i.e., where the prior uncertainty in the load is highest. However, it may be the case that the structural capacity at this location is very high, such that it can withstand any of the potential values the load can take; it is of little use to place the sensor where we are already certain the structure is in no danger. What if, conversely, we were to place the sensor where the load was highest, or where the structure was weakest? Again, if the load was so high or the capacity so low as to guarantee a failure, we would want to apply reinforcements in this area regardless of the load measurement. The context of the larger structural management problem is therefore important to consider when selecting a sensing plan.

What if we instead selected an area where the load was close to the capacity, but still fairly uncertain, such that we were not sure whether the structure would fail or not? This might be a good solution, but if there are multiple areas that fit this criterion, it would be difficult to prioritize between them. One approach might be to prioritize the locations where failure has greater consequences. However, at a certain point, these consequences could grow so high as to be unacceptable, and a choice would be made to reinforce the area regardless of what the loading might be. Economic concerns such as the consequences of failures and the costs of intervention actions also have a role to play in guiding sensor placement.

Now, assume that in one region of the structure the load is known to be quite uniform; a measurement in any part of this region will determine the load applied to the entire region, and therefore allow appropriate actions to be taken throughout. It might therefore be useful to make the measurement here, and thereby to effectively increase the amount of information collected using a single measurement. However, if the consequences of failure in this region are very low, and/or if failures are nearly impossible or virtually certain compared to other parts of the structure, this additional information may be of little relevance. The ability to make wide-ranging inferences based on collected data has the ability to multiply its usefulness, but is not necessarily a guarantee that the resulting information will be worthwhile.

Finally, imagine it were possible to purchase many sensors for the same price as the one being considered, with the consequence that each sensor would have a lower accuracy. Could this increase in data quantity offset the decrease in quality? Could multiple sensors be used to compensate for a lack of precision, while still measuring different areas? The ability to consider the accuracy of collected information, and to trade off different types of potential data collection methods, is another important consideration for sensing scheme design.

The questions raised in this motivating example are typical of the concerns relevant to sensor placement in a wide variety of infrastructure systems. This work will present approaches to answering these questions, and demonstrate them on a variety of specific examples. In particular, Chapter 6 will revisit this motivating example in greater detail, and provide answers to the fundamental questions raised here.

1.2. Background

This review begins with a brief overview of the current practice in management of large-scale distributed infrastructure systems, which tend to focus on management decision-making at a component-level, i.e., by examining individual components in the system, deciding upon optimal actions for these components (potentially including inspection actions), and undertaking these actions without considering their impacts on the larger system. Newer approaches to real-time SHM of infrastructure also tend to have this atomistic view, focusing on assessing individual components without translating this information into system-level health assessments or prioritizing data collection based on system-level infrastructure management needs.

To move from such a component-level infrastructure management approach to decision-making at the system-level, effective system models and the tools to use them for optimal planning are needed. In terms of system models, PGMs and network-based models are outlined here, as they are appropriate to the tasks of modeling uncertainties and interdependencies within infrastructure systems, as discussed above. Finally, existing tools necessary for the use of such models in optimal infrastructure management, including approaches to optimal sensing and applications of decision theory within these models, are briefly presented.

1.2.1. Current Infrastructure Monitoring Practices

Current practice in the management of large-scale infrastructures is based on historical methods of scheduled inspection and maintenance interventions, based on design lifetimes and observed degradation rates of components, supplemented by additional inspections to monitor previously identified potential problems (e.g., Endrenyi et al., 2001). Such methods are usually conservative; inspections are scheduled frequently enough to be able to identify most potential problem conditions before they lead to the failure of the inspected component. However, due to shortages of funding and/or qualified inspectors, these schedules are not always adhered to, e.g., in the case of delayed dam inspections (*Dam Safety*, 2009). Scheduled inspections are also inherently focused at component-level, i.e., are focused on assessing the status of individual components of an infrastructure system, without much interest in quantitatively aggregating this information to assess the status of the system as a whole. Decision-making for infrastructure management is likewise traditionally handled at component-level, with maintenance, repair, or replacement actions scheduled for individual components without much consideration for the system, except in so far as the available budget for these actions must be allocated between many components within the system.

Potential improvements over these traditional monitoring and maintenance practices are often closely related to the way in which information is collected and used to maintain existing infrastructure. A strong positive correlation has been found between proactive maintenance strategies, based on active collection of information and condition-assessment of components, and the performance of these components (Swanson, 2001). These proactive maintenance

strategies contrast with more traditional reactive maintenance strategies in which urgent repairs are performed in response to component failures. This finding illustrates a benefit to the appropriate use of information in supporting infrastructure management in decision-making, and is echoed in many other results. For example, the possible extension of bridge service lives using continuous condition assessment of existing bridges, such as might be provided using state-of-the-art SHM systems, has been conservatively estimated to result in a 31% reduction in bridge replacement costs in the United States (Inaudi, 2011). A reduction of 57% in wind turbine operation and maintenance costs has been estimated for turbines equipped with condition-based monitoring equipment, i.e., suites of sensors that assess the condition of various critical turbine components (Byon and Ding, 2010). A savings of $24 billion over 50 years has been estimated for leak reductions in water distribution pipelines in the UK if instrumented with wireless leak-detection systems (Morimoto, 2010). A threefold increase in the cost-effectiveness of operation and maintenance of railways is also predicted to result from a switch from reactive to proactive railway maintenance (Stenström et al., 2016). These findings illustrate the potential value to society of new monitoring and maintenance strategies based on timely collection and processing of information.

Both traditional monitoring and maintenance approaches, as well as the proposed information-based improvements to these approaches discussed above, conduct the management of entire systems by focusing on the component-level, i.e., they seek to improve system management by improving the management of each component. For example, a common method for proactive maintenance is to develop either probabilistic or deterministic degradation models for individual components, scheduling inspections or maintenance actions for individual components according to these models (Ariaratnam et al., 2001; Kleiner, 2001). In some cases, where the system performance is simply an aggregate of component states and component behaviors are independent, this method will work well. However, in many civil infrastructure systems, there are some components that are objectively more important to the proper functioning of the system than others. These might include components with higher replacement or failure costs (e.g., a major suspension bridge), higher importance to society (e.g., a power substation serving a hospital), or occupying some bottleneck position in the system (e.g., a single pumping station serving multiple water mains). In these cases, component-level degradation models, which ignore interdependencies between components and the relative importance of components within a system, can lead to a sub-optimal allocation of resources. For this reason, a system-level approach to infrastructure monitoring and management is needed.

1.2.2. Structural Health Monitoring

Structural health monitoring (SHM) is a relatively new approach in infrastructure inspection, involving the placement of sensors on or within civil infrastructure components to provide continuous data acquisition for physical parameters of relevance to the performance of the monitored component, e.g., temperature, strain, or acceleration (Farrar, 2013). Most existing and

proposed SHM methods focus on identifying the vibrational modes of structures and inferring damage conditions from this information (e.g., Fan and Qiao, 2011). With this goal in mind, techniques for optimal placement of these sensors on particular structures have been developed, with optimality being measured by the accuracy of identification of modal parameters (e.g., Papadimitriou, 2004; Yi et al., 2011; Yi and Li, 2012; Yuen and Kuok, 2015), or in the reduction of type I and II errors (i.e., false positives and false negatives, respectively) in the identification of structural damage conditions (e.g., Guo et al., 2004; Flynn and Todd, 2010).

A qualitative guideline for the design of SHMs system along these lines is given by Inaudi (2009, 2010), which involves identifying potential risks, uncertainties, and opportunities for monitoring, then defining observable quantities related to these, and finally designing a sensor network to monitor these quantities, considering project lifetime and available budget. However, this approach needs to be further developed into a quantitative framework before it can be widely adopted by infrastructure managers. A current hurdle to the widespread application of SHM technology and techniques in practice is the lack of substantial quantitative assessments of the effectiveness of these systems in terms of reducing life-cycle costs (Inaudi, 2011; Glisic, 2015). Furthermore, for managers of systems of infrastructure, charged with maintaining multiple structures, there is the question of how to optimally deploy new SHM systems within their existing inventories. For many years, budget constraints and institutional inertia will likely delay the deployment of SHM technologies, and so there will be a need for managers of large-scale systems of infrastructure to allocate their resources for new SHM systems among the many components they are tasked with maintaining.

1.2.3. Infrastructure System Modeling

In order to move to a system-level approach to infrastructure monitoring and management, models of interdependent systems must be developed which can describe the overall functioning of these systems and the ways in which their components interact. There are currently several approaches to this problem, which can be grouped into four basic categories: economics-based input-output models, agent-based models, network theory models, and probabilistic models (Satumtira and Dueñas-Osorio, 2010). This overview focuses on the latter two types. The former types, while useful in many contexts, have certain shortcomings that make them less suited to the direction of system-level management. Input-output models, such as those developed by Min et al. (2005), are useful for economic analysis, especially in predicting how disruptions in one sector will ripple through the economy to affect other sectors. However, they are less useful for modeling the initial causes of these disruptions, which are often due to deterioration or extreme events, rather than economic factors. In terms of agent-based models, the computational and programming effort necessary to model large infrastructure systems is still prohibitive in many cases (Satumtira and Dueñas-Osorio, 2010).

1.2.3.1. Probabilistic graphical models for infrastructure

Probabilistic models are a method for quantifying the inherent uncertainty in parameters of physical systems, including infrastructures. By modeling uncertain quantities as random variables, the effects of their variability on the overall system can be explicitly captured and quantified. PGMs are probabilistic models that represent random variables within a graphical visual structure. Such a representation allows the probabilistic relationships between these variables to be easily interpreted. This provides benefits, both in terms of computation (as algorithms have been developed that exploit the graphical structure of the model to perform probabilistic inference efficiently) and in terms of accessibility of the models (as the models are easy to interpret visually, they can be constructed by domain experts with limited knowledge of the underlying probabilistic theories). Finally, in addition to random quantities, deterministic quantities and agent actions, undertaken in accordance with pre-defined strategies, can be incorporated into these models, creating so-called "influence diagrams" or "decision graphs" that can model the actions of rational agents within a stochastic system. PGMs and influence diagrams have thus become an important tool for scientists and engineers seeking to understand the behaviors of systems in which certain basic parameters are inherently uncertain and variable. A thorough introduction to PGMs and influence diagrams is provided by Koller and Friedman (2009).

Probabilistic models have been applied to the analysis of civil infrastructure. Models of the behavior of individual components have been used for reliability analysis, i.e., determining the probability that random variables defining the behavior of these components will fall within certain defined failure regions of the joint probability space (Ditlevsen and Madsen, 1996). For time-varying reliability analysis and for the analysis of systems with multiple components, more advanced methods, incorporating hierarchical PGMs, dynamic Bayesian PGMs, influence diagrams, and structural reliability methods for component-level analysis have been proposed (Straub and Der Kiureghian, 2010a, 2010b). The current state-of-the-art in terms of the modeling of infrastructure systems using PGMs is demonstrated by Bensi et al. (2011, 2015). In that work, a log-Gaussian PGM for seismic risk applied to a proposed system of distributed transportation infrastructure in southern California is presented. This model is then used to guide decision-making for the management of this system. In addition, Luque and Straub (2016) have recently developed a hierarchical dynamic Bayesian system management model. This model captures both the system-level causes and consequences of failures, e.g., in an application to a deteriorating steel truss in which deterioration rates are correlated between members and the capacity of the truss is a function of the number of intact (non-deteriorated) members. Both of these recent works employ discretization of continuous variables defining the system (i.e., the seismic intensity, or the deterioration level) to improve computational tractability.

These latest results for the modeling of infrastructure systems with PGMs also highlight one of the primary difficulties facing the application of these modeling techniques to large-scale infrastructure systems. As the number of components within the system increases, more variables

become necessary to describe the system. These variables occupy a higher-dimensional joint probability space. As the dimensionality of this space grows, traditional methods of exploring the space become less effective. For discrete random variables, the number of possible joint states of the variables grows exponentially, and thus an exhaustive search of this joint space quickly becomes infeasible (e.g., Lincke, 2002). Even approximate techniques are affected; the number of samples required by sampling-based inference algorithms for statistical guarantees on algorithm performance grows rapidly with model dimensionality, to the point where these algorithms are useless without an inordinate number of random samples (e.g., Bertelè and Brioschi, 1972; Koller and Friedman, 2009). This problem falls under the general heading of the "curse of dimensionality", and has been encountered often in high-dimensional data analysis and machine learning (Bellman, 1984).

Methods for lifting this curse are necessary to enable PGM-based infrastructure modeling approaches to scale well to larger systems. One the one hand, careful organization of system-level PGMs to limit the growth of the joint space can allow for increased efficiency in the application of inference algorithms, as has been partially explored by Bensi et al. (2013). On the other hand, approximate methods, such as the use of lower-dimensional approximations to the full joint variable space or simplifying assumptions on the behavior of the system, can be used to avoid some of the issues of these high-dimensional systems.

1.2.3.2. Network models of infrastructure

Another approach to modeling of interdependent infrastructure is based on graph theory and network science. These approaches have a long history of application to a variety of problems, including the modeling of civil infrastructure networks, especially transportation and resource distribution networks. Techniques from graph theory and combinatorial optimization have been applied to solve problems of efficient route planning in transportation networks (Schrijver, 2013).

Network science refers to the analysis of graphical networks, and is related to graph theory (Newman, 2010). Although many applications of network science are in the social sciences, there have been some applications of these techniques to the analysis of infrastructure systems. For example, Buldyrev et al. (2010) and Parandehgheibi and Modiano (2013) use network models to analyze cascading failures in interdependent energy and telecommunications grids, and Dehghani et al. (2014) use network models of transportation networks to assess the vulnerabilities of these networks to the failures of links, in terms of the disruption to network efficiency and increase in travel times. Propagation of failures through networked infrastructure is also discussed by Dueñas-Osorio et al. (2007), who incorporate conditional failure probabilities based on geographic proximities into their analysis to investigate how these stochastic relationships interact with the network structure to affect cascading failures. An example of the state of the art in network disruption analysis (in this case, due to earthquakes) in a realistic infrastructure system (the San Francisco bridge and highway transportation network) is

provided by Baker et al. (2015). Unfortunately, in such realistically-scaled systems (this example considered about 3,000 bridges and 30,000 road links between them), the computational effort required is non-trivial; in this example, only 40 potential seismic scenarios were analyzed, requiring several weeks of time on a high-performance computer cluster.

A key idea of network science, the use of centrality metrics to identify "important" elements in graphical networks, has also been applied to civil infrastructure. This task has been posed as an instance of the "key players problem", i.e., the identification of elements of the network that will, when removed, maximally disrupt the system in some way (Borgatti, 2006). For infrastructure systems, these would be the components that, when they fail, most significantly impact the ability of the network as a whole to provide critical services to society. From this basic idea of what constitutes an "important" component in an infrastructure network, several metrics and heuristics have been developed for identifying these components in infrastructure networks (e.g., Cadini et al., 2009; Kröger, 2011; Galvan and Agarwal, 2015).

The current state of the art in the application of network science techniques to civil infrastructure focuses on modeling the deterministic, rather than probabilistic, relationships between components in these systems. Most models that incorporate uncertainty in component-level performance do so under the assumption that component states are independent of one another (Gertsbakh and Shpungin, 2011). An example of this assumption is given in the work of Liu and Frangopol (2006), who analyze the system-level disruptions caused by component (i.e., bridge) failures in a transportation system. They propose a methodology whereby component-level deterioration models are combined with a system-level assessment of the consequences of component failures, in this case by assessing the user costs of these failures in terms of travel time increases due to the disruption of the system. However, the proposed model, while addressing the relevance of correlations between the behaviors of different components, ignores these correlations for the sake of computational tractability, and instead assumes independent component failures. In practice, common stressors such as climate or increased traffic load can affect many components in the system together, thereby correlating these components' states. Therefore, while addressing the system-level *consequences* of component failures, the proposed model fails to capture the potential system-level *causes* of these failures. Additionally, in the work of Hu et al. (2015), for bridges subject to natural deterioration, the assumption is made that no two bridges will fail simultaneously, and thus bridge failures can be treated as independent events. While this assumption is reasonable for the specific case that they investigate, it will not hold true in all cases, and may lead to serious under-estimation of the system disruption under extreme events (e.g., earthquakes) where multiple simultaneous failures are possible. Unfortunately, abandoning these independence assumptions opens these techniques to the same "curse of dimensionality" discussed for PGMs.

1.2.4. Optimal Sensor Placement

Models of infrastructure systems, no matter what their form, need calibration or input data so that they can provide actionable conclusions about real-world systems. However, the reverse is also true, in that these models can be used to guide the optimal collection of data in the real world by identifying which uncertain parameters have the most influence on the system, such that the reduction of this uncertainty will improve management strategies for the system. This is the essence of optimal sensor placement based on pre-posterior analysis, i.e., using a system model to guide the optimal collection of information such that, when these data are used along with the model, the results will be optimal by the chosen sensor placement objective.

Optimal sensor placement is generally considered a problem in combinatorial optimization, or the selection of some subset of a large potential set, e.g., a small set of sensor placements from a large set of potential sensor placements (Cook et al., 1998; Schrijver, 2013). In contrast to continuous optimization, in which convexity of the objective often allows for efficient solution by convex optimization algorithms, combinatorial optimization is rarely possible in an exact, efficient manner. Many heuristics exist, often inspired by natural processes, such as natural selection (used by genetic algorithms), cognition (neural networks), crystal formation (simulated annealing), and ant colony organization (Colorini et al., 1996). However, for a guaranteed optimal solution, an exhaustive search of the solution space is often necessary. Such searches are computationally prohibitive for all but the smallest problems (e.g., Currin et al., 1991). Therefore, for combinatorial optimization in general and for optimal sensor placement in particular, development of efficient near-optimal solution algorithms is a topic of interest.

Several issues of optimal sensor placement are examined in depth in the seminal work of Krause (2008). A basic solution approach to sensor placement optimization that is presented in that work relies on submodular sensing metrics. Submodularity is intuitively understood as a diminishing returns property, and holds in general for several metrics that might be applied to sensor placement (e.g., Clark et al., 2016). For example, measurements from two sensors might provide a certain amount of redundant information, and therefore the second sensor placed would give less unique information than it would have if operating on its own. In situations where this submodularity property holds, the greedy optimization algorithm, which iteratively adds elements from the potential set of all possible sensors to the set of selected sensor placements, provides provably near-optimal results to combinatorial optimization problems (Nemhauser et al., 1978). Making use of this idea, and the fact that certain information-theoretic metrics such as entropy and mutual information satisfy the conditions of submodularity, Krause presents several approaches to optimal sensor placement. In particular, he focuses on the selection of informative variables to observe in PGMs generally, and in Gaussian process models specifically (Krause et al., 2008c). These techniques are demonstrated in several infrastructure-relevant applications, such as the detection of contamination in a water supply system (Krause et al., 2008a), lighting control optimization using sensors for building energy usage reduction (Singhvi et al., 2005), and determining temperature and precipitation patterns (Guestrin et al., 2005). These results

demonstrate the developed techniques for generating provably near-optimal sensor networks for the monitoring of systems modeled with Gaussian process models. Recent work has extended these techniques to optimizing sensor placement for detecting areas where values (again modeled by Gaussian random fields) may exceed a certain threshold, or optimize a value function over a domain (Gotovos et al., 2013; Sui et al., 2015; Vanchinathan et al., 2015). These methods have potential applications to the placement of sensors for detecting regions of concern for distributed infrastructure subject to spatially varying risks.

1.2.5. Value of Information

In terms of the practical application of sensor placement techniques to infrastructure management, both in the case of state classification as discussed for SHM systems in Section 1.2.2 and the conditional entropy metric in Section 1.2.4, considerations relating to the economics and management of the system are not accounted for. The importance of such considerations in infrastructure management is illustrated by an intuitive thought experiment. Consider a system of two structures, for which the resources are only available to inspect one. A prior probabilistic performance model of these structures estimates their probabilities of failure at 50% for one and 1% for the other. Under either the classification error reduction approach or the conditional entropy metric, an inspection should be performed for the former structure, since its state is the most uncertain, i.e., its state would only be correctly predicted 50% of the time, as opposed to 99% of the time for the other component. However, for many practical applications in infrastructure, a 50% probability of failure is unacceptably high, because the expected cost of component failure would be much higher than the cost of repairing or replacing the component to prevent a failure. Therefore, in a practical application, the first component (with a high probability of failure) would be replaced immediately, while the second component (with a low probability of failure) might be inspected to determine if a replacement activity is truly necessary. This example emphasizes the importance of considering potential management actions and their respective costs when making decisions about infrastructure inspections. Neither the conditional entropy nor misclassification rate metrics take these factors into account, and therefore may not be practically applicable to the optimal monitoring of infrastructure. An alternative metric that does explicitly account for these factors is the decision-theoretic value of information (VoI).

Decision theory refers to the development of mathematics- and statistics-based approaches to rational decision-making under uncertainty. These approaches often take the form of strategies that will, by some metric, maximize the long-term expected utility to the decision-maker of their chosen actions, given a range of uncertain outcomes. In relation to infrastructure management, ideas from decision theory can be used to guide optimal maintenance strategies that will reduce the expected life-cycle costs of the infrastructure, taking uncertainties in its performance into account. One approach to applying decision-theoretic techniques to the management of infrastructure under uncertainty involves incorporating manager actions and user outcomes into

system PGMs, creating influence diagrams that can be used to analyze potential management strategies. Approaches that model infrastructure systems using PGMs often adopt this approach to optimize decision-making for these systems (e.g., Bensi et al., 2011).

It should be noted that other approaches to decision-making under uncertainty exist which have goals besides the maximization of expected utility. These include approaches where the worst-case performance is to be maximized, or where a certain percentile of the distribution of potential outcomes (e.g., the outcome with a 95% probability of exceedance by the utility measure) is to be maximized. There are also approaches which seek to optimize performance over the range of possible outcomes, although this can be analyzed in a decision theoretic framework by considering all outcomes to be equally likely and applying standard expected utility maximization (e.g., Lempert and Collins, 2007). Satisficing approaches also exist, which seek to find a solution which, while not optimal by the utility measure, satisfies minimum performance criteria (Simon, 1959). Still other strategies seek to maximize flexibility as the best approach to decision-making under uncertainty in sequential decision-making problems (Rosenhead, 2013). These various robust decision-making strategies are typically applied in cases of deep uncertainty, where the probability distribution over potential outcomes cannot be defined to the decision-maker's satisfaction. In this work, it is assumed that potential outcomes, their consequences, and their probabilities will be known to the decision-maker, and therefore the basic expected utility decision-making framework will be adopted (robust approaches are briefly revisited in a discussion of sensor placement under model uncertainty in Section 11.3).

The design of optimal information-collection strategies for infrastructure, such as inspection schemes or SHM systems, can be examined from a decision-theoretic perspective. A key metric for this analysis is the VoI, defined as the difference in expected costs for managing a system without and with collecting certain additional information (Raiffa and Schlaifer, 1961). Computation of the VoI is a form of pre-posterior analysis, in that estimation of the VoI is made before the information itself is collected, and therefore an expected value of the outcome must be taken over the probability distribution for the possible information that might be observed. A complete mathematical introduction to VoI is deferred until Chapter 4, with a discussion of the computational issues associated with this metric presented in Chapter 5. Here, a more qualitative overview of the metric and its application is presented.

Since its formulation, the VoI concept has been applied to many diverse fields, including accounting (e.g., Hilton, 1977) and medicine (e.g., Heckerman et al., 1989; Claxton and Posnett, 1996). The use of the VoI metric also has clear applications in the assessment and design of sensing systems and inspection strategies for infrastructure, due to its potential to quantify the utility of information gathering efforts, and therefore help to determine whether such efforts will be cost-effective. An important early work in the application of VoI in civil engineering was that of Straub (2004), which developed strategies for reliability-based inspection planning of deteriorating components, using decision-theoretic concepts to identify the timing of inspections that would result in the lowest life-cycle management cost for the component. Already at this

stage, the difficulties in optimizing inspections at a system-level were identified; the exact application of VoI to system-level inspection scheduling was considered computationally prohibitive, and an approximate heuristic approach was adopted based on iterative optimization of schedules for individual components (Straub and Faber, 2004, 2006).

More recent work has expanded on the applications of VoI in civil infrastructure. Assessing the long-term VoI provided by SHM systems is discussed by Pozzi and Der Kiureghian (2011a, 2011b). Applications of the VoI metric to supporting decision-making for geotechnical investigations of dikes and levees has also been an area of significant interest (e.g., Qin et al., 2015; Schweckendiek and Vrouwenvelder, 2015; Thöns et al., 2015). The use of the VoI as a heuristic for allocating inspection efforts between components of a system is discussed by Memarzadeh and Pozzi (2016).

Recent computational advances have also made VoI analysis possible in larger problems, but many barriers to its application at system-level remain. For example, a sampling-based approach to calculating the VoI of inspections on a deteriorating component is outlined by Straub (2014); this approach is efficient only for relatively small-scale systems (i.e., single components of an infrastructure system), as the sampling approaches presented scale poorly with increasing numbers of random variables. Additionally, optimal inspection of infrastructure is examined by Goulet et al. (2015a, 2015b); sensing activities and structural interventions for the monitoring and maintenance of a pair of similar columns are optimized by cost, subject to meeting a required safety level (i.e., a maximum prescribed probability of failure). Again, this approach scales poorly with the size of the system, as an exhaustive enumeration of all potential management strategies for the pair of columns is required.

One potential method for addressing the poor scaling performance of VoI-based sensing optimization is to adopt an approach similar to that of Krause (2008), where a greedy optimization heuristic is used to iteratively add potential observations of the system based on maximizing the VoI until a specified “budget” for information collection activities is exhausted. Unfortunately, as shown by Krause (2008), VoI is not generally a submodular metric, and therefore the same guarantees on solution optimality cannot be made. However, this lack of submodularity does not necessarily invalidate the greedy approach to optimization; it may still be of interest to pursue this approach, even without rigorous guarantees on optimality, for the sake of increased computational efficiency and scalability. Such an approach is therefore adopted in this work, and some of its consequences are examined in Section 11.1.

Recent work has aimed to identify submodular surrogates for the VoI based on the decision region determination problem (Chen et al., 2015). In this problem, the space of possible hypotheses, i.e., possible discrete states for the system, is divided between decision regions, i.e., subsets of the hypothesis space under which a given action is optimal, using the smallest number of binary tests which distinguish true and false sets of hypotheses. In that work, an adaptive submodular surrogate metric is developed which, while aligning with the goals of the VoI, allows for efficient and near-optimal selection of informative measurements. Although

promising, the approach as presented is limited by the need to consider discrete system states and test outcomes. Furthermore, generalizing the approach to apply to imperfect tests (i.e., noisy measurements) involves augmenting the hypothesis space with new hypotheses corresponding to different noisy outcomes, which exacerbates the exponential state space growth problem. For these reasons, such an approach is not adopted in this work.

In a recent application to optimal planning of geotechnical investigations, Yoshida (2015) uses a Gaussian process model of arsenic concentration in soil, as well as a VoI metric based on the relative costs of type-I and type-II errors in classifying the arsenic concentration levels and a particle swarm optimization technique, to plan investigations of sites for characterization of soil contamination. This approach allows for sensor placements to be optimized in the continuous domain of the random field, rather than among a set pre-selected sensor placements on a discretized grid spanning the domain. Therefore, some problems of sub-optimal performance of the greedy algorithm are avoided, but at an increased computational cost, as the gradients of the VoI field at the location of each sensor placement must now be computed, in addition to the VoI quantities themselves. Additionally, the particle swarm optimization approach can still find itself trapped in local optima, as there is no evidence that the problem of sensor placement optimization via the VoI metric is convex.

In summary, while the VoI metric has applications to assessment and design of monitoring systems or inspection schemes for civil infrastructure, its application in large-scale systems (with many components with correlated behaviors, for which a variety of potential measurements and management actions are possible) is hindered by the computational effort necessary, as compared with other sensor metrics (e.g., conditional entropy). Methods for the efficient computation and optimization of this metric, either exactly or through incorporating appropriate simplifying assumptions in specific systems or problems, are therefore a potentially fruitful area of research.

1.2.6. Summary of Research Gaps

When examining the body of prior work on infrastructure system monitoring and management, several recurring themes are present that suggest gaps in the current research:

Component-level thinking: decisions on monitoring and management in systems are typically made for components without considering the system-level impacts of these decisions. Even many proposals for new SHM systems consider “big-data” approaches, i.e., using vast numbers of cheaply deployed sensors to monitor every component in a system, rather than “smart-data” approaches (Sheth, 2013), i.e., focusing sensing efforts where information will have the most benefit for managing the system as a whole. Transitioning from a component-level to a system-level outlook entails major computational challenges (discussed under the next point), but also has potentially large benefits in terms of more efficiently utilizing limited resources.

Scaling issues: the analysis of infrastructure systems becomes more difficult as the number of components in the system grows. This issue arises as a result of system-level coupling between components, which can come in several forms, such as system-level causes (i.e., a common influencing factor affecting an entire system at once, such as in the case of an extreme event like an earthquake), system-level effects (i.e., the state of a system depending in a complicated way on the states of its components, such as in transportation network routing and connectivity), system-level information (i.e., when information gathered within a system comes from several of its components, such as in cases where spatially averaged data are collected), and system-level actions (i.e., management activities affecting more than one component of the system). Recent computational improvements have made steps towards alleviating these difficulties, but the fundamental challenge of the "curse of dimensionality" remains.

Value-based sensing: accounting for the economic implications of collected information on decision-making, as is accomplished explicitly by VoI, is important to assessing how useful information is in supporting infrastructure management. An information-theoretic approach, focusing on uncertainty reduction alone, may not capture all of the relevant issues. A more thorough investigation of where information-theoretic and decision-theoretic approaches both agree and disagree in terms of optimizing sensor placement is therefore important to identify situations in which either or both approaches can be applied appropriately.

1.3. Research Questions

Within the context of the prior work outlined above, this work examines how existing concepts from probability and statistics, i.e., the use of PGMs and influence diagrams to describe decision-making for uncertain systems, together with concepts for optimal sensor placement and decision theory, i.e., the VoI, can be applied to support the management of infrastructure systems. The fundamental motivating assumption for this research is that infrastructure monitoring and management will be most efficient and effective when both tasks are considered together, i.e., when decisions are made optimally based on collected information, and when this information is collected optimally to support decision-making. Following from this assumption, the general research question to be addressed in this work is:

> **How can probabilistic models of infrastructure systems, together with the value of information metric, be used to guide the collection of information to optimally support the management of these systems?**

The use of the term "optimally" in the research question is meant to reflect concepts of both computational and economic efficiency, i.e., any proposed solution should quantify the optimality of sensing based on the underlying economics of infrastructure management, and do so in a manner that is computationally tractable for the system considered.

This general research question is addressed throughout this work by way of examinations into the following six specific research sub-questions:

1. What is an appropriate model for an infrastructure system that can capture interdependencies in component performance, observation activities, management actions, and their consequences, and thus support the optimization of sensing and management together?
2. Under what assumptions can the VoI be used to support optimal sensor placement for the management of infrastructure systems in a computationally efficient manner?
3. For situations where the VoI cannot be efficiently computed, what alternative or approximate techniques might be applied to optimizing sensor placement?
4. What is the relationship between VoI and other sensor placement metrics, in terms of how different factors relevant to infrastructure management applications are accounted for, i.e., how is "optimality" in sensing defined by different metrics, with respect to infrastructure systems?
5. Can VoI be efficiently applied to placement and scheduling of sensing in evolving systems, and under what assumptions can this efficient application be made?
6. How can the practical application of the VoI metric for supporting infrastructure system monitoring and management be demonstrated?

1.4. Key Results and Contributions

This work presents several results that contribute to answering the general research question posed above. Several important contributions are tied directly to the specific research sub-questions:

1. Chapter 3 describes the PGM and influence diagram framework for infrastructure systems adopted in this work. This intuitive, flexible framework is used to describe the relationships between infrastructure systems, their constituent components, observations of their states and properties, actions taken to manage these systems, and their consequences. The general framework is further used to identify systems where the VoI metric can be efficiently computed.
2. Using the probabilistic modeling framework described above, specific cases in which VoI computation can be performed efficiently are identified. Based on these cases, an efficient computational approach is presented for calculating VoI in systems with a specific topology (dubbed the "cumulative" topology) under certain assumptions on the system (i.e., that the variables describing system performance are jointly Gaussian, observations of these are also Gaussian, component states are binary, and management actions have local effects on specific components). This computational approach is described in Chapter 5.
3. Cases in which the assumptions described above are not met by the problem structure are also investigated. In Section 5.3, a generally applicable approximate method of VoI computation is outlined, along with certain efficient techniques applicable to series and parallel system topologies in particular. Furthermore, in Section 5.4, an investigation into

the performance of the VoI metric for cases in which the system does not exhibit the cumulative topology are investigated numerically for computationally tractable systems, and compared with results for corresponding cumulative systems, in an effort to draw parallels between these cases that can form the basis of approximate methods.

4. Theoretical comparisons of the VoI with other sensor placement metrics are provided in Section 4.4. Furthermore, several applications of these metrics to example sensor placement problems are presented in Chapter 6, where the different characteristics of these metrics, in terms of the optimal sensor placements resulting from different metrics under different circumstances, are compared.
5. Chapter 7 presents an extension of the VoI computation and sensor placement methodologies discussed previously to evolving systems, i.e., systems where the underlying factors influencing the performance of the infrastructure system are time-dependent. A temporal variant of the cumulative topology assumption is presented and applied to support efficient computation of the VoI metric in evolving systems. This metric can then be used to support both optimal placement and scheduling of information collection in space and time.
6. Along with many theoretical demonstrations of the application of the VoI metric, three larger specific case studies of how the VoI metric can be employed for optimal sensor placement to support infrastructure management are presented in this work. In Chapter 8, a set of strain gauges installed on the Sherman and Joyce Bowie Scott Hall are used as a testbed to demonstrate the application of the optimization techniques developed in this work to supporting infrastructure management decision-making for a simulated problem of differential settlement between column foundations. In Chapter 9, the problem of seismic risk in the San Francisco Bay area is examined, and an optimal scheme for monitoring and managing a set of bridge and tunnel infrastructures in the region in the aftermath of an earthquake is described. In Chapter 10, the problem of extreme heat risk in an urban area is considered, and a probabilistic model of urban temperatures is developed and used to support the optimal placement of new weather stations in the city of Pittsburgh, Pennsylvania to support heat wave advisory issuance. Together, these case studies provide specific demonstrations of how probabilistic models and the VoI metric can be used to support the management of diverse infrastructures, ranging from individual structures (i.e., the Scott Hall building), to portfolios of physical structures (i.e., bridges and tunnels in the San Francisco Bay area), to social infrastructures supporting warning issuance and emergency response activities (i.e., heat warning issuance for the city of Pittsburgh).

1.5. Overview of the Document

This document is organized into 12 chapters, including this introductory chapter. The outline for subsequent chapters is as follows:

Chapter 2 introduces the general problem of sensor placement optimization for infrastructure system management. It introduces the notation for random fields, component and system states, management actions and losses, and observations that is used throughout the work. It also discusses some general issues for sensing in infrastructure, and describes the greedy approach to sensor placement optimization which is adopted in this work.

Chapter 3 outlines the PGMs and influence diagrams used to model infrastructure systems in this work. It presents various system topology effects that are incorporated into the general model, and discusses the characteristics of the cumulative system topology in particular. It gives a general background on the Gaussian random field modeling approach used throughout the work. Finally, it presents the assumptions of binary component state definition and linear observations that are used to define states and observations in this work.

Chapter 4 introduces the VoI metric, providing definitions and important properties of the metric which have been identified in previous work. A parametric analysis of the metric for a single-component system is presented to provide insight into how this metric captures properties of interest for infrastructure monitoring efforts. Finally, the metric is compared with other sensor placement metrics, such as posterior predictive error and conditional entropy, to illustrate equivalences and differences between these.

Chapter 5 discusses the computational challenges associated with evaluating the VoI metric in large systems. A general approximate method for VoI computation is presented. Furthermore, an efficient exact method for computation under the assumption of a cumulative system topology is presented. A comparative example is analyzed for the VoI in cumulative, series, and parallel system topologies, illustrating the behavior of the metric in these different multi-component systems.

Chapter 6 gives comparative examples for sensor placement based on the VoI and other metrics in one-dimensional and two-dimensional problems. The effects of system topologies, state definitions, possible actions, and alternative metrics are investigated. These examples serve to illustrate how the VoI captures relevant features for supporting infrastructure management decision-making in various contexts.

Chapter 7 extends the methodologies developed in this work to problems in a spatio-temporal domain, i.e., where random fields vary not only across the system in space, but also over time for a finite discretized time horizon. Efficient evaluation of the VoI in such systems is illustrated based on the cumulative topology assumption discussed previously. This extension allows for the application of this metric to optimal sensor scheduling as well as placement.

Chapters 8 through 10 illustrate demonstrative applications of the techniques developed in this work to specific infrastructure management problems. First, in Chapter 8, a recent project to instrument the columns of the Sherman and Joyce Bowie Scott Hall on the Carnegie Mellon University campus is used as a demonstrative case for optimal sensor placement to support SHM. Next, in Chapter 9, a Poisson process model of seismic risk in the San Francisco Bay area

is used to support optimal sensing for post-earthquake management of bridges and tunnels. Finally, in Chapter 10, a spatio-temporal Gaussian process model of extreme heat risk is developed and calibrated, and used along with the VoI metric to support temperature monitoring for optimal heat wave advisory issuance for the city of Pittsburgh.

Chapter 11 discusses open issues related to the use of the VoI metric to support optimal sensor placement for infrastructure management. These include the lack of submodularity of the VoI metric, and its impact on the potential performance of the greedy optimization algorithm. Some anecdotal results and empirical evaluations of algorithm performance are presented. Other open issues include the evaluation and optimization of VoI under model uncertainty and the efficient analysis of large-scale systems with non-cumulative topologies via this metric.

Finally, Chapter 12 summarizes the results of this work, and provides some general conclusions about the use of PGMs and VoI to support optimal sensor placement for infrastructure system monitoring and management. It also lists several publications in which the primary results presented in this work first appeared.

Chapter 2

Problem Statement

This chapter defines key terms and concepts that form the foundation of this work. These include the definition of random fields affecting infrastructure systems, the definition of the states of these systems and their components, of observations taken within the system, and of actions taken for managing the system and their consequences, all of which are discussed in Section 2.1. Furthermore, the general concept of optimal sensor placement within such a system is presented in Section 2.2, and the greedy optimization approach which will be used throughout the work is introduced in this context.

Before beginning with mathematical definitions, a brief qualitative introduction to the concepts of infrastructure systems and their monitoring and management, as these ideas are understood and used in this work, is in order. The term “infrastructure system” is used generally to describe any engineered structure or structures that function to provide a basic service for society. Examples range from individual buildings, to groups of buildings or structures, to networks of many diverse constructs. A transportation system will be used in the following paragraphs as a running example to help the reader visualize the concepts discussed.

Infrastructure systems are made up of components. In a transportation system, these components might be bridges, tunnels, and roadways. In this work, the distinction between system and component is arbitrary, based on the specific problem or issue to be solved, and no attempt will be made to provide a general taxonomy for systems and components. Instead, such distinctions will arise naturally from context, and may be fluid between different applications; for example, in one context a bridge many be considered a system, with its beams and connections as components, and in another the bridge many be the component, with the system made up of numerous bridges connected into the transportation network. The one distinction that will be adopted generally between components and systems is their spatial extent; the components of infrastructure systems are considered to be distributed across a spatial domain, where each component occupies a specific point within that domain. However, this is again a contextual rather than an absolute distinction, as in reality all components are three-dimensional objects, and thus are themselves distributed across space. In context, however, the difference in scale between a system and its components often makes it convenient to describe each component as being associated with a discrete point within the larger domain of the system.

The term “state” is used to refer to the operational condition of an infrastructure component or system. For the system, this state refers to its ability to or the degree to which it can fulfill its designated societal function. For a transportation system, this state could refer to the connectivity between two cities, or to the traffic capacity the network is able to sustain between them. For a component, this state refers to its ability to or the degree to which it can fulfill its designated function as a part of the system. A bridge’s state within a transportation network might then refer to whether or not it is open for traffic, or what traffic loading it can sustain. In this work, it is assumed that the state of a system can be described via a deterministic or probabilistic function of the states of its components.

The behaviors of all system components, and therefore the behavior of the system, are uncertain. In other words, an agent acting to manage the system has incomplete and/or imperfect knowledge of the states of the components and of the various factors that determine these states. For example, an agency managing a transportation system does not have complete knowledge of the condition states of its bridges, the environmental conditions to which they are subjected, or the traffic loadings applied to them. An agent can, however, build a mathematical model of the system that takes such uncertainty into account. In Bayesian parlance, this knowledge and uncertainty is captured by a prior probabilistic model of the system. This uncertainty can be reduced through sensing, which is broadly defined as any activity aimed at reducing uncertainty in the system. This can be understood traditionally, as the placement of physical sensor within the system to measure variables of interest, or more generally, as the performance of inspections or the use of improved modeling techniques to reduce uncertainty. Uncertainty reduction is accomplished, in a Bayesian sense, by using the collected information to update the prior probabilistic model of the system to a posterior model, reflecting this new information. Monitoring of the system is defined as the set of sensing efforts made by the managing agent to collect information within the system to reduce his or her uncertainty in its performance. The problem of “sensor placement” in this work therefore refers to the choice of what information collection efforts to make within the system.

The managing agent of the system is not a passive observer of the system only, but also has the option to actively intervene. For example, a transportation system manager has the option to replace damaged bridges, retrofit existing bridges, or restrict access to bridges that may be unsafe. The management of the system is defined as the set of active interventions in the system taken by the managing agent. Note that this set must include the empty set, or the “do nothing” activity, i.e., the active choice by the managing agent not to intervene in the system. The choice of what actions to take in managing an uncertain system must be guided by the ultimate goals of the agent and of society. In this work, a decision-theoretic or utilitarian outlook is adopted. It is assumed that the managing agent has a utility function, i.e., a mapping from the joint space of all variables affecting the system and all actions undertaken for its management to a scalar quantity ranking this outcome’s desirability with respect to all other possible outcomes. Typically, this utility function is assessed in a monetary form and in the inverse, i.e., as a “loss function” that quantifies the disutility or cost of an outcome to the managing agent on a monetary scale. This is

due to the desire to connect sensing and decision-making to the economic considerations of infrastructure management, which is accomplished most conveniently by using such a scale. Furthermore, adoption of the pessimistic viewpoint (i.e., seeking to minimize losses rather than to maximize gains) is arbitrary, based on the preference of the agent; adopting the optimistic viewpoint would be completely equivalent, and it is left to the reader to translate between these viewpoints as desired.

Taking this view of system monitoring and management, the problem of optimal sensor placement can be addressed quantitatively. Specific quantification of the problem using the VoI metric is deferred until Chapter 4. For now, the problem is described generally as the selection of a set of information collection activities that maximizes some objective mapping this set to a quantitative assessment of the relative merits of this set with respect to other potential sets. It is taken for granted that that sensing quality can be quantified in such a way; again, a justification and comparative discussion of different methods of performing this quantification is differed to Chapter 4. This chapter will discuss the problem generally, as well as the solution approach adopted here.

The following sections give a more concrete mathematical formulation to the concepts described qualitatively above.

2.1. Basic Definitions

To begin, $\mathcal{X}$ defines the spatial domain over which the system is distributed. This might be a one-, two-, or three-dimensional space; furthermore, higher dimensions may be considered in the case that components are co-located, such that they might be distinguished. Furthermore, in some cases, a non-spatial coordinate system may be used, e.g., if a more abstract system of social infrastructure is being modeled, features such as age, sex, education, etc. might be relevant to localize the "components" of this system within a social space. However, for the vast majority of cases considered in this work, $\mathcal{X}$ denotes a standard Euclidian spatial domain.

For computational purposes, this continuous spatial domain is discretized into a finite set of n_X locations, collected into the set X. This is sometimes used simply to represent the continuous space $\mathcal{X}$ through a finite number of points on a regular grid. However, it is often the case in this work that this discretization corresponds to the locations of the components within the system, as only these locations will be of direct relevance to system management. In such cases, the number of discrete locations n_X matches the number of system components n. It is assumed that these locations are known, i.e., the coordinates of each component of the system are fixed quantities and the components do not move.

2.1.1. Random Fields

The many factors influencing the behavior of the components of the system are described using random fields defined over the domain $\mathcal{X}$. Denote by F the set of random fields affecting the

system, and denote by $\mathrm{f}(\mathbf{x})$ the random field values affecting the system at location $\mathbf{x} \in \mathcal{X}$. In general, $\mathrm{f}(\mathbf{x})$ may be multivariate, describing the values of multiple co-located random fields for different features of interest in modeling the system. Furthermore, variables may be modeled that are not associated with a specific location $\mathbf{x} \in \mathcal{X}$, but apply generally throughout the system; in such cases, for consistency, these variables can be associated with an arbitrary "dummy" location, or multiple identical copies of the variable may be associated with each $\mathbf{x} \in X$.

The vector $\mathbf{f}$ denotes the random quantities that affect the system at all of the discretized locations X. In the case where there are multiple co-located random fields for different quantities, the values of all fields are concatenated into this single vector for notational simplicity.

The potential values of random vector $\mathbf{f}$ are described according to its prior probability distribution:

$$\mathbf{f} \sim \mathrm{p}_F \tag{2-1}$$

The random variables affecting system performance may in general be continuous or discrete, depending on the context. Although their values are unknown, a key assumption for most of this work is that these variables are static, i.e., that they take on a fixed (although unknown) quantity that does not change. In the context of infrastructure system management, this assumption can be viewed in three ways. First, these variables can represent physical properties that are deterministic but unobserved, i.e., the consequence of epistemic uncertainty. A specific example of this might be the ultimate strength of a beam, which is a fixed quantity but is unknown before the beam's failure. Second, these variables can represent realizations of a random process during a specific extreme event, i.e., the consequence of aleatory uncertainty. An example of this might be the peak loading on a structure during a hurricane, which may vary from one hurricane to the next but is fixed within a given hurricane event. Third, these variables can represent long-term averages or trends. For example, while the loading on soil can fluctuate greatly, the average loading over time may define the soil consolidation (Das, 2010, chap. 11). In Chapter 7, the problem of dynamic systems which evolve over time is considered; in that case, different instantiations of each variable in $\mathbf{f}$ are considered for different times over the distretized temporal domain. Therefore, again, the subset of variables corresponding to each time is static, with past and future instantiations of these quantities represented by different variables, rather than the same variables change over time.

2.1.1.1. Model parameters and scenarios

It is further assumed in this work that any model parameters W, i.e., any parameters of the probabilistic model defining the prior distribution, are known. For example, in a multivariate Gaussian distribution, these parameters refer to the mean vector and covariance matrix of the distribution. Equation 2-1 can then be understood as the distribution of $\mathbf{f}$ conditional to known model parameters $\mathbf{w}$:

$$\mathbf{f}|\mathbf{w} \sim \mathrm{p}_{F|\mathbf{w}} \tag{2-2}$$

In general, these parameters may be unknown, described instead by prior distribution p_W. In this case, the prior distribution for **f** can be derived using the law of total probability:

$$\mathrm{p}_F(\mathbf{f}) = \int_{\mathbf{w}\in\mathrm{Dom}(W)} \mathrm{p}_{F|\mathbf{w}}(\mathbf{f}|\mathbf{w})\, \mathrm{p}_W(\mathbf{w})\, \mathrm{d}\mathbf{w} \tag{2-3}$$

In Chapter 9, this assumption is investigated further, where parameters **w** are used to describe the earthquake scenario affecting a system in terms of its magnitude and epicenter location. In that case, the seismic scenario affecting the system is itself uncertain, and therefore uncertainty in **w** is accounted for. However, elsewhere in this work, unless explicitly mentioned, the assumption of fixed model parameters, i.e., of a known scenario, is followed.

2.1.2. Observations

Sensing activities can be modeled as taking measurements of the underlying random field F. A set of such observations is denoted Y, and the distribution of the measured values **y** of these observations conditional to the underlying random field variables **f** is:

$$\mathbf{y}|\mathbf{f} \sim \mathrm{p}_{Y|\mathbf{f}} \tag{2-4}$$

The prior distribution for these measurements can be defined using the law of total probability:

$$\mathrm{p}_Y(\mathbf{y}) = \int_{\mathbf{f}\in\mathrm{Dom}(F)} \mathrm{p}_{Y|\mathbf{f}}(\mathbf{y}|\mathbf{f})\, \mathrm{p}_F(\mathbf{f})\, \mathrm{d}\mathbf{f} \tag{2-5}$$

In cases where these measurements are associated with specific points in the domain, $\mathrm{y}(\mathbf{x})$ denotes a measurement of $\mathrm{f}(\mathbf{x})$. In general, these measurements may not be associated with specific points, e.g., these may be measures of sets or functions of random field values for one or many points within the domain. Also, these measurements are generally not precise, i.e., $\mathrm{y}(\mathbf{x}) \neq \mathrm{f}(\mathbf{x})$, but may have noise or systematic bias. In any case, the conditional distribution $\mathrm{p}_{Y|\mathbf{f}}$ defines the relationship between the measurements and the random field quantities.

2.1.2.1. Bayesian updating

Through Bayesian inference, measurements **y** are used to define a posterior distribution for the random field variables **f** conditioned on this observation. This posterior distribution is:

$$\mathbf{f}|\mathbf{y} \sim \mathrm{p}_{F|\mathbf{y}} \tag{2-6}$$

where this posterior distribution is defined via Bayes' theorem:

$$\mathrm{p}_{F|\mathbf{y}}(\mathbf{f}|\mathbf{y}) = \frac{\mathrm{p}_{Y|\mathbf{f}}(\mathbf{y}|\mathbf{f})\, \mathrm{p}_F(\mathbf{f})}{\mathrm{p}_Y(\mathbf{y})} \tag{2-7}$$

The posterior distribution descibes the updated understanding of the underlying random variables affecting the system gained as a result of the information $\mathbf{y}$ gathered from the observation set Y. In general, the definition of such a posterior distribution can be difficult, requiring approximate techniques such as Markov Chain Monte Carlo (MCMC), which allows samples to be drawn from this posterior distribution without an explicit definition of its form (e.g., Au and Beck, 2001; Koller and Friedman, 2009; Papaioannou et al., 2015). However, in certain special cases a so-called "conjugate prior" distribution is used, and the inference can be performed in closed form. One such case is the multivariate Gaussian distribution, which is its own conjugate prior. In Chapter 3, the use of the Gaussian distribution for modeling infrastructures is discussed.

2.1.3. States

The definitions of the states of infrastructure components and systems are presented here, in which three different state variable types are presented. *Limit-state variables* are intermediate variables between the random field variables affecting the components and c*omponent state variables*, and *system state variables* describe how the system as a whole performs. In this section only the prior distributions for each of these types of state variables are described. However, posterior distributions for these can easily be derived using the relationships between these variables and the random field variables presented in this section and the posterior distribution for the random field variables described in Equation 2-7. Furthermore, note that the relationships between these variables are described probabilistically. However, in Chapter 3 and beyond, these relationships are assumed to be deterministic; such deterministic relationships may be understood as special cases of the general probabilistic description provided here.

2.1.3.1. Limit-states

Limit-state variables are used here, as by, e.g., Ditlevsen and Madsen (1996), to provide an intermediate step between the underlying random variables affecting the system and the states of the system components. For example, when multiple co-located random fields are affecting the same component in different ways, it is often useful to transform these multiple variables into a single variable describing the overall effect of these variables on the component. The set of limit-state variables for the system is denoted G, and the distribution for $\mathbf{g}$, the vector of variable values, conditional to the random field variables $\mathbf{f}$, is:

$$\mathbf{g}|\mathbf{f} \sim \mathrm{p}_{G|\mathbf{f}} \tag{2-8}$$

The prior distribution for $\mathbf{g}$ is obtained from the law of total probability:

$$\mathrm{p}_G(\mathbf{g}) = \int_{\mathbf{f}\in\mathrm{Dom}(F)} \mathrm{p}_{G|\mathbf{f}}(\mathbf{g}|\mathbf{f})\, \mathrm{p}_F(\mathbf{f})\, \mathrm{d}\mathbf{f} \tag{2-9}$$

Limit-state variables are primarily used for notational convenience, as they describe the influences of multiple factors using a single variable, and allow for easier definition of

component states (these often being dependent on a single limit-state variable rather than on multiple random field variables). However, the use of limit-state variables is not necessary in all cases, and these are then ignored, i.e., a definition of $\mathbf{g} = \mathbf{f}$ is adopted.

2.1.3.2. Component states

The set of component states S describes the behavior of each of the n components of the system. Vector $\mathbf{s}$ encodes these states, and its distribution conditional to the limit-state variables $\mathbf{g}$ is:

$$\mathbf{s}|\mathbf{g} \sim \mathrm{p}_{S|\mathbf{g}} \tag{2-10}$$

The prior distribution for $\mathbf{s}$ can be obtained using the law of total probability:

$$\mathrm{p}_S(\mathbf{s}) = \int_{\mathbf{g}\in\mathrm{Dom}(G)} \mathrm{p}_{S|\mathbf{g}}(\mathbf{s}|\mathbf{g})\, \mathrm{p}_G(\mathbf{g})\, \mathrm{d}\mathbf{g} \tag{2-11}$$

Component state variables describe the level of functionality of each component of a system. This level may be descibred by a continuous or discrete variable, e.g., a binary (yes or no) indication of functionality, a bridge condition rating (an integer from 0 to 9), or the residual capacity in a structural element (a real number).

In general, multiple state variables may be associated with each component, but in this work one variable is assigned to each component, i.e., it is assumed there are n variables in vector $\mathbf{s}$. Because of this, there may be different ways to interpret the concept of "component" in different contexts. Most commonly, components are considered to be separate physical entities that have a certain functional behaviour. However, the same physical entity may be divided into several different components for the purpose of system analysis, e.g., if the physical entity in question has multiple operational or failure modes, or fulfills several purposes within the system. In such cases, each of these virtual components is associated with a state, rather than each physical component.

2.1.3.3. System state

The state of the system defines its functionality with respect to the service the system is designed to provide for society. The system state variable $\mathbb{s}$ is defined conditional to the component states as:

$$\mathbb{s}|\mathbf{s} \sim \mathrm{p}_{\mathbb{S}|\mathbf{s}} \tag{2-12}$$

This variable can be discrete or continuous, as for the components above. The prior distribution of the system state is defined using the law of total probability:

$$\mathrm{p}_{\mathbb{S}}(\mathbb{s}) = \int_{\mathbf{s}\in\mathrm{Dom}(S)} \mathrm{p}_{\mathbb{S}|\mathbf{s}}(\mathbb{s}|\mathbf{s})\, \mathrm{p}_S(\mathbf{s})\, \mathrm{d}\mathbf{s} \tag{2-13}$$

The method whereby the component states contribute to the system state can be physical (e.g., a network of bridges providing transportation connectivity) or logical (a structure with several failure modes, each of which constitutes a virtual component contributing to the state of the system). Furthermore, the system state may not be defined as a single variable for all system topologies, but rather via the collection of component state variables, i.e., $\mathbb{s} = \mathbf{s}$ for some systems. Such cases will be examined further in Chapter 3.

2.1.4. Actions

Actions denote the activities undertaken by an agent managing the system. Based on his or her prior knowledge of the system, as well as any information collected to update this prior knowledge, the managing agent chooses a set of actions A from a set of potential actions $\mathcal{A}$. The choice of management actions is encoded into vector $\mathbf{a}$. In contrast to the terms introduced previously, $\mathbf{a}$ is not a random variable. Rather, it is a decision variable, whose value is chosen by the managing agent. However, in practice, assuming a rational decision-maker and a known loss function, an optimal decision can be defined conditional to a set of measurements or random field values. However, the precise manner in which the choice of action is decided is left unspecified until Chapter 4.

Two key assumptions are made in defining actions and their consequences in this work. First, it is assumed that actions are non-sequential, i.e., that all actions are decided upon and executed simultaneously throughout the system. This connects to the assumption of a static random field stated in Section 2.1.1; as there is no temporal distinction between the underlying random variables, there is no need for any temporal distinction between management actions. In other words, the problem of sequential decision-making, in which the proper order of actions must be determined, is not considered for the majority of this paper. An exception is in Chapter 7, where this problem is considered under specific assumptions on the structure of the sequential decision-making problem.

Second, it is assumed that actions cannot influence the random field values, and therefore the states of the components. This may seem counter-intuitive, e.g., “how can a failed component be repaired if actions cannot affect the component state?” However, in this work these effects are captured by the loss function, as discussed below. This method for describing the effects of actions is important to prevent the formation of cycles in the PGM of the system which would lead to computational difficulties, as discussed in Chapter 3.

2.1.5. Losses

The loss function is a map from the space of random variables affecting the system and the space of actions taken for system management to a scalar value that quantifies the desirability of the managing outcome to the agent, with a higher value corresponding to a less desirable outcome. In its most general form, the loss function is denoted $\mathrm{L}(\mathbf{f}, \mathbf{a})$, a function of random field variables

$\mathbf{f}$ and management actions $\mathbf{a}$. In certain contexts, the loss function may also be formulated as $\mathrm{L}(\mathbf{g},\mathbf{a})$, $\mathrm{L}(\mathbf{s},\mathbf{a})$, or $\mathrm{L}(\mathbb{s},\mathbf{a})$, a function of limit-state variables $\mathbf{g}$, state variables $\mathbf{s}$, or system state variable $\mathbb{s}$ and actions $\mathbf{a}$ respectively. These various formulations can be based on the relationships outlined in Section 2.1.3. For example:

$$\mathrm{L}(\mathbf{g},\mathbf{a}) = \mathbb{E}_{F|\mathbf{g}}\mathrm{L}(\mathbf{f},\mathbf{a}) = \int_{\mathbf{f}\in\mathrm{Dom}(F)} \mathrm{L}(\mathbf{f},\mathbf{a})\ \frac{\mathrm{p}_{G|\mathbf{f}}(\mathbf{g}|\mathbf{f})\ \mathrm{p}_F(\mathbf{f})}{\mathrm{p}_G(\mathbf{g})}\ \mathrm{d}\mathbf{f} \tag{2-14}$$

and in the reverse direction:

$$\mathrm{L}(\mathbf{f},\mathbf{a}) = \mathbb{E}_{G|\mathbf{f}}\mathrm{L}(\mathbf{g},\mathbf{a}) = \int_{\mathbf{g}\in\mathrm{Dom}(G)} \mathrm{L}(\mathbf{g},\mathbf{a})\ \mathrm{p}_{G|\mathbf{f}}(\mathbf{g}|\mathbf{f})\ \mathrm{d}\mathbf{g} \tag{2-15}$$

The appropriate formulation of the loss function depends on the decision-making problem context, with it being more convenient in some cases to define the loss in terms of $\mathbb{s}$, $\mathbf{s}$, or $\mathbf{g}$ rather than $\mathbf{f}$. Some common loss function formulations for different situations will be discussed in Chapter 4.

The definition of the loss function as depending on both variables and agent actions allows the effects of these actions to be captured by the loss function directly. This also allows the costs of these actions to the managing agent to be captured by the loss function, i.e., both the cost of implementing an action and the result of this action in affecting the performance of the system are captured in the loss function. This formulation, however, requires that the outcomes of the actions be deterministic. In other words, given an action and the state of the system (described by $\mathbb{s}$, $\mathbf{s}$, $\mathbf{g}$, or $\mathbf{f}$ depending on the loss function), there is a known outcome for the action. In principle, actions with uncertain outcomes, e.g., repairs that may or may not succeed, can also be modeled in this framework by augmenting the state vector, e.g., instead of "failed" and "operational" states, there may be "failed", "failed, but reparable", and "irreparably failed" states, with the outcome of the "repair" action depending on which of the latter two states holds. However, the use of an augmented state in this manner is not investigated in this work.

2.2. Optimal Sensor Placement

The problem of optimal sensor placement is formally stated as:

$$Y^* = \mathrm{argmax}_{Y\in\mathcal{Y}}\ \mathrm{M}(Y) \quad \text{subject to} \quad \mathrm{C}(Y) \le b \tag{2-16}$$

where the optimal set of measurements Y^* is selected from the set of all possible measurements $\mathcal{Y}$ to maximize the sensor placement objective $\mathrm{M}(Y)$ subject to the constraint that the cost of obtaining these measures $\mathrm{C}(Y)$ is less than or equal to the sensing budget b.

The definition of the sensing objective $\mathrm{M}(Y)$ is left unspecified for now. For the majority of this work, the VoI metric will be used to define $\mathrm{M}(Y)$, as discussed in Chapter 4. However, alternative metrics, such as misclassification rate (which is in fact a special case of the VoI

metric) and the conditional entropy, will also be employed. For now, only a qualitative discussion of the properties of the objective is provided.

Sensing objective $\mathrm{M}(Y)$ is a set function of measurement set Y that maps this set to a scalar quantity. This quantity is a measure of the quality of the information provided by measurement set Y, where a higher quantity indicates a higher quality of information. Note that it is not a function of the measurement outcomes $\mathbf{y}$ directly, but only of the measurement scheme Y. The objective should therefore be robust against the prior uncertainty in the outcome of these measurements, i.e., it should be defined using a pre-posterior analysis of the potential quality of the measurement set across all possible measurement outcomes. Beyond this, the way in which the quality of the measures is evaluated is left open, i.e., no specific dependence on F, G, S, A, etc. is specified.

Broadly speaking, there are two potential approaches to defining this objective that are examined in this work. First, there is the information-focused approach that has as its motivation the idea that the ultimate goal of sensing is to reduce uncertainty in the system. Following this approach, the sensing objective should be defined in such a way that it measures how the uncertainty in the random variables defining the system will be reduced following an observation of Y. The conditional entropy and misclassification rate metrics follow directly from this approach. Second, there is the outcome-focused approach, which is motivated by the idea that the goal of sensing is ultimately to allow for better decision-making in managing the system, i.e., to reduce losses or maximize the utility of the decision-making agent in the uncertain system. This approach is the one adopted in this work, and the VoI metric discussed in Chapter 4 follows from it. However, these two approaches are not always mutually exclusive, and there are certain similarities in how both approaches are used to define sensor placement objectives. These similarities are investigated in Section 4.4 for the metrics discussed.

2.2.1. Optimization Approach

Obtaining a solution to Equation 2-16 is an exercise in combinatorial optimization. As discussed in Section 1.2.4, exact solution of a combinatorial optimization problem is in general intractable. For example, let $n_{\mathcal{Y}} = |\mathcal{Y}|$, the number of individual candidate measurements in set $\mathcal{Y}$. The number of possible subsets (including the empty set) of $\mathcal{Y}$ that might be chosen is then $2^{n_{\mathcal{Y}}}$. Due to this exponential growth, an exhaustive enumeration and evaluation of all potential candidate measurement sets will not be feasible, except for problems where $n_{\mathcal{Y}}$ is small (the limit such that an exhaustive search can be executed depends on the computational effort and time required to test each candidate measurement set). Therefore, an approximate solution approach is needed to efficiently solve this combinatorial optimization problem in general.

The approach adopted in this work is the application of the greedy optimization algorithm. This algorithm works to build the putative optimal set of measurements by at each step of the algorithm adding the single remaining unselected measurement in $\mathcal{Y}$ that maximizes the sensing

objective when added to the set of measurements already selected in previous steps. Therefore, instead of simultaneously selecting all measurements in the proposed set, only one measurement at a time is selected and added to the set. This approach requires at most $\frac{1}{2}n_{\mathcal{Y}}\left(n_{\mathcal{Y}}+1\right)+1$ evaluations of the objective function to evaluate potential measurement sets, with fewer evaluations needed if the sensing budget is exhausted before all $n_{\mathcal{Y}}$ candidates are included in the selected set. Algorithm 1 presents pseudo-code for the greedy optimization algorithm as applied in this work. This code is adapted from similar algorithms presented by Krause (2008). Note that $\cup$ denotes the set union and $\backslash$ the set subtraction. Although there is in general no guarantee on the performance of this algorithm, its simplicity and efficiency recommend it to the task of obtaining solutions to Equation 2-16 for large systems. Investigations into the performance of this algorithm for the problems investigated in this work are provided throughout, and a summary of its performance is given in Section 11.1.

Algorithm 1: Pseudo-code for the greedy algorithm.

Input candidate set $\mathcal{Y}$, objective function $\mathrm{M}(\cdot)$, cost function $\mathrm{C}(\cdot)$, constraint b
$j=0,\ \mathcal{Y}_0=\emptyset$
for each $Y\in\mathcal{Y}$
    if $\mathrm{C}(Y)\leq b$ **then** $\mathcal{Y}_0=\mathcal{Y}_0\cup\{Y\}$
end
while $\left|\mathcal{Y}_j\right|>0$
    $j=j+1$
    select Y_j^* (See Equations 2-17 and 2-18)
    $\mathcal{Y}_j=\mathcal{Y}_{j-1}\backslash\left\{Y_j^*\right\}$
    for each $Y\in\mathcal{Y}_j$
        if $\mathrm{C}\left(\left\{Y_1^*,\dots,Y_j^*,Y\right\}\right)>b$ **then** $\mathcal{Y}_j=\mathcal{Y}_j\backslash\{Y\}$
    end
end
$j_{\mathrm{end}}=j$
Output $Y^*=\left\{Y_1^*,\dots,Y_{j_{end}}^*\right\}$

Note that the algorithm presented above gives an example of a forward greedy optimization algorithm, i.e., the optimal set is built one element at a time by adding elements to an initially empty set. It is also possible to adopt a reverse greedy optimization algorithm where, beginning with set $\mathcal{Y}$, elements are removed from this set until an optimized subset is achieved (e.g., Papadimitriou, 2004). However, this approach will require more computational effort in typical scenarios where $n_{\mathcal{Y}}$ is large but the number of sensors in the optimal set is small, as many iterations of the algorithm must be performed, each including a large number of observed variables, before the set can be trimmed to an appropriate size. Section 7.2.3 provides an example of a reverse greedy optimization approach being implemented, and Section 11.1 further examines such approaches.

In Algorithm 1, the method for selecting the optimal candidate measurement Y_j^* at each iteration depends on whether sensor placement is being performed in an offline or online manner.

2.2.1.1. Offline Optimization

In offline sensor placement, all measurements are chosen simultaneously, such that no observations are available before the full measurement set has been chosen. Under this placement regime, the next proposed measurement to be added to the optimal set is:

$$Y_j^* = \text{argmax}_{\,Y_j \in \mathcal{Y}_{j-1}} \mathrm{M}\left(\left\{Y_1^*, \dots, Y_{j-1}^*, Y_j\right\}\right) \tag{2-17}$$

That is, the next selected measurement is that which maximizes the sensor placement objective when combined with all previously selected measurements. Offline optimization is the most common approach considered in this work. As previously mentioned, random field variables are presumed to be static. Therefore, their measurements will also be static. Furthermore, it is assumed that all management actions are undertaken simultaneously. Therefore the order in which measurements are obtained is typically not relevant to the problem application.

2.2.1.2. Online Optimization

In the case of online sensor placement, measures are selected sequentially, with observations from all previously selected measurements being available to guide the choice of the next measurement in the sequence. This is also known as adaptive optimization, sequential optimization, or closed-loop optimization of sensor placement. In this case, the next proposed measurement to be added to the optimal set is:

$$Y_j^* = \text{argmax}_{\,Y_j \in \mathcal{Y}_{j-1}} \mathrm{M}\left(Y_j | \mathbf{y}_{j-1}\right) \tag{2-18}$$

where $\mathrm{M}\left(Y_j | \mathbf{y}_{j-1}\right)$ denotes the sensor placement objective evaluated for Y_j where previously selected measures $\{Y_1^*, \dots, Y_{j-1}^*\}$ have as their outcomes the observation vector $\mathbf{y}_{j-1}$. In this case, because observations of $\left\{Y_1^*, \dots, Y_{j-1}^*\right\}$ are already available, there is no need to take into account uncertainty in these previous measurements; only the outcome of Y_j is still uncertain. Online optimization is therefore more computationally efficient in general; rather than requiring pre-posterior analysis for all selected measurements, only uncertainty in the most recently selected measure must be accounted for. Furthermore, for any given system, an online sensor placement will perform the same as or better than an offline sensor placement in terms of the final sensing objective value obtained. The reason for this is the inclusion of the additional measurement information in guiding the placement. Intuitively, by including this information, a posterior model is obtained for the system that more accurately captures its characteristics, and therefore the resulting sensor placement guided by this posterior model will be better suited to the system than a measurement guided by the prior model. A further discussion of this in the context of the

VoI metric is provided in Section 4.1.3. Finally, note it is still assumed that all measurements are taken and used to update the probabilistic model before any final management actions are implemented, i.e., there is no sequential decision-making.

Although online sensor placement is not commonly employed in this work, an example using the VoI metric is considered in Section 6.2.5. Furthermore, in Chapter 7, for the case of sensor placement and scheduling in spatio-temporal systems using the VoI metric, there is an inherent sequence to the measurement activities, and therefore online placement is a natural approach.

Chapter 3

Modeling Infrastructure Systems

In this chapter, the definitions and relationships introduced in Chapter 2 are used to create a PGM and influence diagram of an infrastructure system. A general model is introduced in Section 3.1, and its key characteristics discussed. In Section 3.2, several special cases of this general model that are of interest for this work are presented. Finally, a set of system modeling assumptions used throughout this work are introduced in Section 3.3. These are:

1. Random fields that affect the system can be modeled using a Gaussian process framework; equivalently, probabilistic random field model p_F of Equation 2-1 is a multivariate Gaussian distribution. See Section 3.3.1 for details.
2. Observations of the random field variables can be expressed as linear combinations of random field variables and Gaussian measurement noise, i.e., the conditional distribution $p_{Y|\mathbf{f}}$ of Equation 2-4 is also multivariate Gaussian. See Section 3.3.2 for details.
3. Limit-state variables are described by an affine transformation of the random field variables, i.e., $p_{G|\mathbf{f}}$ of Equation 2-8 is a deterministic relationship, defining a multivariate Gaussian distribution for **g**. See Section 3.3.3 for details.
4. States variables of each component of the system and of the system as a whole are binary random variables, i.e., p_S of Equation 2-11 is a multivariate Bernoulli distribution and $p_{\mathbb{S}}$ of Equation 2-13 is a univariate Bernoulli distribution. See Sections 3.3.4 and 3.3.5 for details.

These assumptions and the corresponding relationships presented in this chapter will be commonly used throughout this work; cases in which they are relaxed will be specifically indicated by the statement “basic modeling assumptions of Section 3.3 are relaxed in this section”.

3.1. Probabilistic Graphical Modelling Approach

The conditional probability relationships presented in Section 2.1 can be used to define a generic PGM of the variables for random fields, observations, and states. Furthermore, this PGM can be expanded into an influence diagram via the inclusion of actions and losses within the model. Figure 3-1 presents the generic influence diagram for an infrastructure system and its management based on the relationships of Section 2.1.

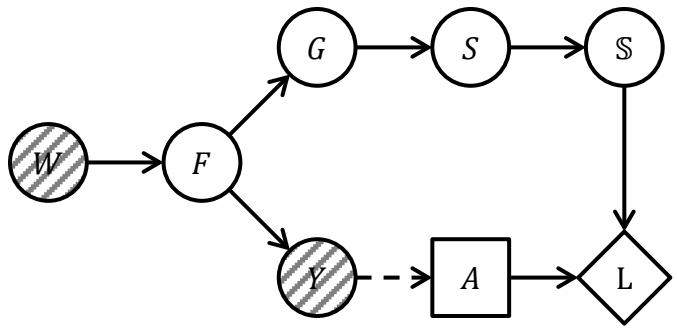

Figure 3-1: Generic influence diagram of an infrastructure system and its management.

The layout of Figure 3-1 follows a common notational convention for influence diagrams (Koller and Friedman, 2009). Circles denote random variables or sets of variables, with shaded circles denoting that the values of these variables are observed. Following the discussion of Section 2.1.1.1, model parameters W are observed, along with measurements Y. Rhombi denote deterministic functions of variables, in this case the loss function L. Squares denote decisions whose values are set by an agent within the system; in this case, the action set A. Arrows between random variables denote the probabilistic dependence structure; e.g., a conditional probability distribution for G conditioned on a value of F is represented by an arrow from F to G. Arrows connecting to deterministic or decision nodes denote functional dependencies, e.g., arrows from $\mathbb{S}$ and A to L denote that the loss is a function of the actions and the system state. Note that this assumes that a loss function of the form $\mathrm{L}(\mathbb{s}, \mathbf{a})$ is being modeled; for other formulations of the loss function discussed in Section 2.1.5, intervening variables are removed from the model. Figure 3-2 presents a simplified influence diagram in which a loss function of the form $\mathrm{L}(\mathbf{f}, \mathbf{a})$ is modeled, which is obtained from the generic model of Figure 3-1 by removing nodes G, S, and $\mathbb{S}$ and connecting F directly with L. Finally, the dashed arrow connecting Y to A does not denote a probabilistic or deterministic relationship, as the agent's choice of action is left free. However, there is a precedence relationship between these, in that the outcomes of the observations of Y will be used by the agent to help guide his or her choice of actions A towards a desireable outcome for the management of the system.

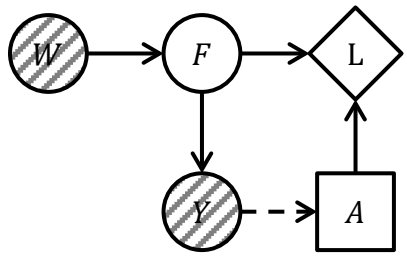

Figure 3-2: Simplified influence diagram, modeling a loss function of the form L(**f**,**a**).

An important fact to note in examining the graphical structure of Figure 3-1 is the lack of cycles formed by the directed edges. This property characterizes the PGM as a Bayesian Network. Such acyclicity of the influence diagram is essential to ensure tractability for the decision-making problem captured in the graphical model (Koller and Friedman, 2009). As alluded to in Section 2.1.4, this is a key reason why it is assumed that actions cannot influence the random field F, as this would introduce a cycle connecting the random field, observations, and actions.

3.2. Infrastructure System Topology

In this work, infrastructure system topology is understood to encompass the various factors that describe how the system is arranged, i.e., how the components interact with one another, and how they are monitored and managed as part of the system. This general concept of topology encompasses four related concepts. The *functional topology* of the system describes how the various random field variables affect specific components of the system in different ways. This topology can be understood as influencing the definitions of $\mathrm{p}_{G|\mathbf{f}}$ and $\mathrm{p}_{S|\mathbf{g}}$ from Equations 2-8 and 2-10. The *network topology* of the system describes how the system behaves as a result of its components' states. This can be understood as defining $\mathrm{p}_{\mathbb{S}|\mathbf{s}}$ from Equation 2-12. The *monitoring topology* describes how measurements are collected in the system. This influences both the definition of $\mathrm{p}_{Y|\mathbf{f}}$ from Equation 2-4 and the set $\mathcal{Y}$ of possible measurements that may be considered in the system. The *management topology* describes how the system can be managed. This identifies both the set $\mathcal{A}$ of possible management actions that can be considered, as well as influences how the loss function is defined. These concepts of topology are related to each other, e.g., network topology has an effect on management topology since the consequences of a component failure should be quantified by the loss function in terms of how this failure influences the operational state of the system through the network structure.

The influence diagram of Figure 3-1 represents a generic model, applicable broadly to any type of infrastructure system topology that can be described using the definitions and relationships presented in Section 2.1. However, based on the specific topology of the system being modeled, alternative influence diagrams representing variations on this generic diagram may be better suited to the system. Such alternatives are motivated by either a desire to trim unnecessary components of the model (such as in Figure 3-2, where nodes that do not influence the loss function are removed) or a desire to add additional details that better reflect the system's characteristics. Three specific alternatives are presented in this section, following this latter motivation, for three classes of system topologies of interest in this work.

3.2.1. Independent Systems

In an independent system, the monitoring and management of all system components are undertaken separately. Figure 3-3 presents an influence diagram for this type of system. The model uses a template representation; nodes enclosed in the dotted box are repeated for each of

the n components of the system. Nodes within this box denote subsets of variables associated with each component, e.g., F_i denotes a subset of F that affects the ith component of the system.

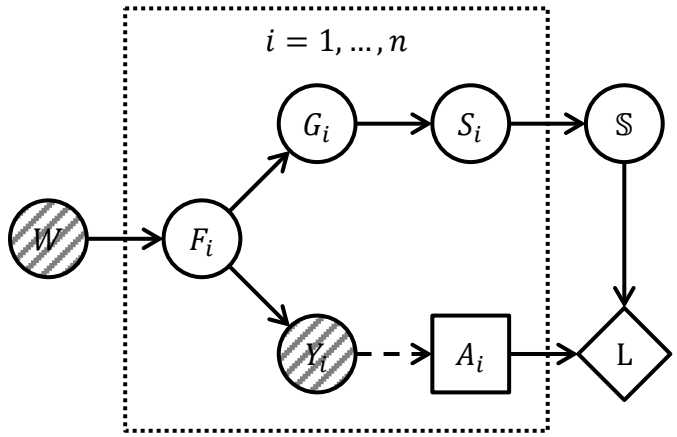

Figure 3-3: Influence diagram for a system with independent components monitored and managed independently.

In this system, the sets of variables affecting each of the system components are independent, conditional to the model parameters W. As a result, the states of these components are also independent. It is also assumed that the observation set Y is divided into subsets, where Y_i denotes a set of observations relating to factors affecting the ith component. As these factors are independent across components, only measurements in the set Y_i are relevant to supporting the management of the ith component. Furthermore, it is assumed that the set of actions can be separated, so that a subset A_i of these actions affects the ith component.

The structure of this influence diagram incorporates several assumptions about the functional topology of the system (i.e., independence between variable sets affecting different components), the monitoring topology (measurements are made locally for each component, which is a consequence of the independence of the random field variables between components), and the management topology (separate action sets are considered for each component). The network topology, i.e., the definition of the system state in terms of the component states, is left unspecified.

This model is relevant in that it captures the current practice for infrastructure system management, as discussed in Chapter 1. Models are typically developed for different system components in isolation; this is reflected by the template structure repeated across the system components, and the independence of variables affecting the components is a consequence of this approach. Monitoring activities and management actions are also typically undertaken at the component-level. System-level interactions are in some cases captured through a system network model. As this model represents the typical current approach to system modeling, it is not

investigated further in this work; it is merely presented and discussed here to provide a point of reference for the other model structures introduced below.

3.2.2. Cumulative Systems

As an alternative to the independent system model, this work will focus heavily on the following model, which is denoted as the *cumulative system* model. An influence diagram of a cumulative system is presented in Figure 3-4. In contrast with the independent system model of Figure 3-3, in the cumulative system the random field variables, limit-state variables, and observations may be dependent across components. Only the component states themselves are independent conditional to the limit-state variables. The assumption of separate sets of management actions for each component is maintained in the cumulative system. Furthermore, a set of component-level loss functions are introduced, such that L_i is a function only of the component state S_i and the management action set A_i for the ith component. The system-level loss function L is then assumed to be the sum of the component-level loss functions, hence the "cumulative" nomenclature for this system.

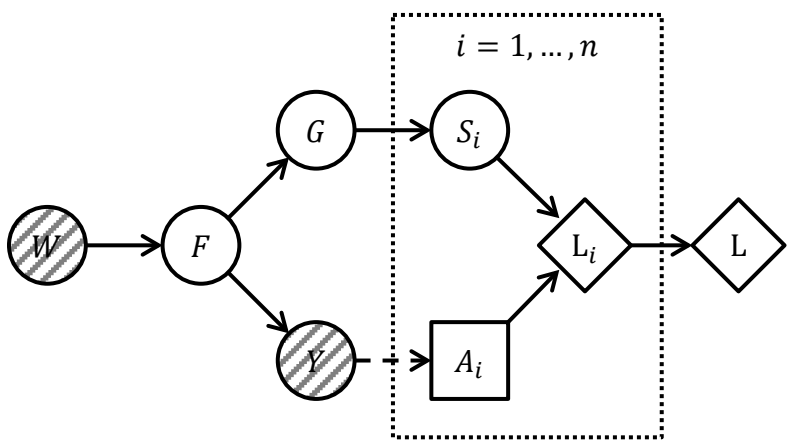

Figure 3-4: Influence diagram for a system with a cumulative topology.

The key assumption for defining a cumulative system is the decomposability of the loss function. For a system to be cumulative, the system-level loss function must decompose across the system components as follows:

$$\mathrm{L}(\mathbf{s}, \mathbf{a}) = \sum_{i=1}^{n} \mathrm{L}_i(\mathbf{s}_i, \mathbf{a}_i) \tag{3-1}$$

Note that this decomposition can also distinguish between the components of the system; for any system-level loss function that is decomposable as in Equation 3-1, the decomposition defines the components of the system, with state $\mathbf{s}_i$ and actions $\mathbf{a}_i$ assigned to the ith component, with the number of components of the system equal to the number of components into which the loss function decomposes. Note that component state and action vectors may be multivariate.

Furthermore, note that loss functions of the form $\mathrm{L}(\mathbf{f}, \mathbf{a})$ and $\mathrm{L}(\mathbf{g}, \mathbf{a})$ may also be decomposable, e.g., as $\mathrm{L}(\mathbf{f}, \mathbf{a}) = \sum_{i=1}^{n} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i)$ or $\mathrm{L}(\mathbf{g}, \mathbf{a}) = \sum_{i=1}^{n} \mathrm{L}_i(\mathbf{g}_i, \mathbf{a}_i)$. Loss functions of the form $\mathrm{L}(\mathbb{s}, \mathbf{a})$ are by definition non-decomposable; such a decomposition would imply that $\mathbb{s} = \mathbf{s}$, i.e., that any definition of $\mathbb{s}$ as distinct from the vector $\mathbf{s}$ is superfluous, and therefore $\mathrm{L}(\mathbb{s}, \mathbf{a})$ would be expressed directly as $\mathrm{L}(\mathbf{s}, \mathbf{a})$. Finally, note that decomposability of the loss function does not imply marginal independence of the component states, limit-states, or random field variables. For instance, in the case of a decomposable loss function $\mathrm{L}(\mathbf{f}, \mathbf{a})$, $\mathbf{f}_i$ and $\mathbf{f}_j$ may still be dependent for $i \neq j$.

The cumulative system model incorporates no assumptions on monitoring or functional topologies, as limit-states G and measurements Y are still defined by joint distributions in the PGM. The cumulative system does assume certain network and management topologies, as encoded in the decomposable loss function. The cumulative system can be considered as the "null" network topology, i.e., the components do not act together to define a system state, other than as the vector of all of the component states, as mentioned above. The cumulative system topology and the implied decomposability of the loss function have important implications for the computational efficiency of sensor placement objective evaluations that are discussed in Chapter 5.

3.2.3. Non-Cumulative Systems

A non-cumulative system influence diagram is presented in Figure 3-5, to contrast with the cumulative system depicted in Figure 3-4. All assumptions of the cumulative system are preserved, except for the decomposability of the loss function across components. In other words, the system-level loss must be expressed as a function of all component states and management actions together.

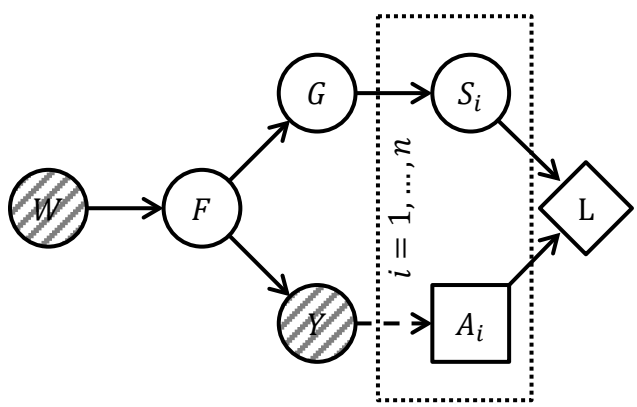

Figure 3-5: Influence diagram for a system with a non-cumulative topology.

The non-cumulative system model presented here is used in this work to examine systems in which the network topology is such that decomposition of the loss function as in Equation 3-1 is

not possible. Two prototypical examples of such systems are those with *series* or *parallel* network topologies. In a series system, the system as a whole will only function if all of its components are functioning, and in a parallel system, the system as a whole will function unless all of its components cease to function. In both cases, the consequences of the system state are nonlinear in the number of operating components, and therefore the loss function will not decompose. These system topologies and their consequences will be examined in further detail elsewhere in this work, and the influence diagram of Figure 3-5 will be used to model such systems.

3.3. Gaussian Random Field Infrastructure Model

This section introduces a set of assumptions on the probabilistic model structure that will be used to support much of the remaining work, as outlined at the beginning of this chapter. The key assumption introduced here is that the set of random fields F affecting the system can be modeled using a Gaussian process modeling framework. The Gaussian process is a generalization of the multivariate Gaussian distribution to a continuous domain (Rasmussen and Williams, 2006). Specifically, random field F defined over domain $\mathcal{X}$ can be described by a Gaussian process model if and only if, for any finite subset of locations $X \subseteq \mathcal{X}$, the joint distribution for the random field values at these coordinates can be described by a multivariate Gaussian distribution. The Gaussian process notation for such a random field is:

$$\mathrm{f}(\mathbf{x}) \sim \mathcal{GP}\left(\mu(\mathbf{x}), \mathrm{k}\left(\mathbf{x}_i, \mathbf{x}_j\right)\right) \tag{3-2}$$

where this model is parameterized by mean function $\mu(\mathbf{x})$, denoting the mean of the marginal Gaussian distribution for $\mathrm{f}(\mathbf{x})$, and covariance function $\mathrm{k}\left(\mathbf{x}_i, \mathbf{x}_j\right)$, which defines the covariance between the random field variables at coordinates $\mathbf{x}_i$ and $\mathbf{x}_j$. The covariance function is decomposable as:

$$\mathrm{k}\left(\mathbf{x}_i, \mathbf{x}_j\right) = \sigma(\mathbf{x}_i)\sigma\left(\mathbf{x}_j\right)\rho\left(\mathbf{x}_i, \mathbf{x}_j\right) \tag{3-3}$$

where $\sigma(\mathbf{x}_i)$ is the marginal standard deviation of $\mathrm{f}(\mathbf{x}_i)$ and $\rho\left(\mathbf{x}_i, \mathbf{x}_j\right)$ defines the correlation coefficient between $\mathrm{f}(\mathbf{x}_i)$ and $\mathrm{f}\left(\mathbf{x}_j\right)$.

Gaussian process models have a long record of use for the modeling of spatially distributed quantities; historically, the use of such models was known as “kriging” in the field of geostatistics (Cressie and Wikle, 2011). Many applications and extensions of the basic Gaussian process framework have been employed in various domains. Of particular interest, the optimal placement of sensors using Gaussian process models and the conditional entropy (or mutual information) metric is investigated by Krause et al. (2008d) and is a primary motivation for the application of Gaussian processes in this work. For further information on Gaussian processes, the reader is referred to the definitive text of Rasmussen and Williams (2006).

For the applications of Gaussian processes presented in this work, several facts should be noted. First, in this context, the set of model parameters W denotes both information on the parametric forms of $\mu(\mathbf{x})$ and $k(\mathbf{x}_i, \mathbf{x}_j)$, as well as their respective parameters. That is, the assumption that the parameter vector $\mathbf{w}$ is known is equivalent to an assumption that the forms and parameters of the mean and covariance functions defining the random field are known. Second, the case of multiple co-located random fields can easily be handled within the Gaussian process framework by augmenting the coordinate vector $\mathbf{x}$ with an indexing value corresponding to the various co-located fields. These fields can then be described using a common joint Gaussian distribution, e.g., by assuming independence between variables relating to different random fields (where such an assumption is appropriate). The values for variables in all fields are then concatenated into the common random field vector $\mathbf{f}$.

Finally, note that this framework, although explicitly defining jointly Gaussian random variables, can also be used to describe other joint distributions so long as an appropriate transformation can be defined between the multivariate Gaussian and the target distribution. For example, a set of log-normally distributed random variables can be described by a Gaussian process model by first applying the logarithm transformation to these variables, and then defining a Gaussian process model for the random field of these variables in the log-space. Such a technique is used in Chapter 9, where the seismic risk field affecting a system is naturally modeled by a jointly log-normal distribution.

3.3.1. Gaussian Random Fields

As discussed at the start of the section, the random field variables affecting the system can be modeled using a Gaussian process framework. The result is a multivariate Gaussian distribution for these variables:

$$\mathbf{f} \sim \mathcal{N}(\boldsymbol{\mu}_F, \boldsymbol{\Sigma}_F) \tag{3-4}$$

where the mean vector $\boldsymbol{\mu}_F = \mu(X)$ is obtained by evaluating the mean function at every spatial coordinate in the set X, and the covariance matrix $\boldsymbol{\Sigma}_F = k(X, X)$ is obtained by evaluating the covariance function at all pairs of coordinates in X. This serves to define the prior distribution of Equation 2-1 under the Gaussian random field assumption. Note that conditioning on model parameters as in Equation 2-2 is implicit through the definition of the mean and covariance functions.

3.3.2. Linear Observations

It is further assumed that observations of random field variables can be described as linear combinations of these variables and a measurement noise, as:

$$\mathbf{y} = \boldsymbol{\Omega}_Y \mathbf{f} + \boldsymbol{\epsilon} \tag{3-5}$$

where $\mathbf{\Omega}_Y$ is the measurement matrix, specifying what variables or linear combinations of variables in the random field are measured. The size of $\mathbf{\Omega}_Y$ is n_Y by n_X, where n_X is the number of elements in X, i.e., the length of random field vector $\mathbf{f}$, and n_Y is the number of measurements in Y, i.e., the length of observation vector $\mathbf{y}$. The measurement noise vector $\boldsymbol{\epsilon}$ is assumed to have a multivariate Gaussian distribution as well:

$$\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{\mu}_\epsilon, \mathbf{\Sigma}_\epsilon) \tag{3-6}$$

where the measurement noise has mean $\boldsymbol{\mu}_\epsilon$ and covariance $\mathbf{\Sigma}_\epsilon$. In this context, the assumption of zero-mean measurement noise would correspond to a choice of $\boldsymbol{\mu}_\epsilon = \mathbf{0}$ (the zero vector), and the assumption of uncorrelated or white noise would correspond to the choice of $\mathbf{\Sigma}_\epsilon$ as a diagonal matrix. Note however that these assumption are not included here in general; measurements may be biased, with correlated noise between different measurements.

As a linear combination of Gaussian random variables, the vector of observations can itself be described by a joint Gaussian distribution. Under the assumptions of this section, the prior distribution of Equation 2-5 becomes:

$$\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}_Y, \mathbf{\Sigma}_Y) \tag{3-7}$$

where:

$$\boldsymbol{\mu}_Y = \mathbf{\Omega}_Y \boldsymbol{\mu}_F + \boldsymbol{\mu}_\epsilon \quad \mathbf{\Sigma}_Y = \mathbf{\Omega}_Y \mathbf{\Sigma}_F \mathbf{\Omega}_Y^{\mathrm{T}} + \mathbf{\Sigma}_\epsilon \tag{3-8}$$

An important property of Gaussian distributions, and a primary motivation for their use in this work, is the fact that a Gaussian distribution is its own conjugate prior. This allows for a posterior Gaussian distribution for $\mathbf{f}$ to be obtained from its Gaussian prior and Gaussian measurements $\mathbf{y}$. In this case, the posterior distribution of Equation 2-6 is:

$$\mathbf{f}|\mathbf{y} \sim \mathcal{N}\left(\boldsymbol{\mu}_{F|\mathbf{y}}, \mathbf{\Sigma}_{F|Y}\right) \tag{3-9}$$

Furthermore, the posterior mean and covariance can be evaluated in closed-form as:

$$\boldsymbol{\mu}_{F|\mathbf{y}} = \boldsymbol{\mu}_F + \mathbf{\Sigma}_{FY} \mathbf{\Sigma}_Y{}^{-1} (\mathbf{y} - \boldsymbol{\mu}_Y) \quad \mathbf{\Sigma}_{F|Y} = \mathbf{\Sigma}_F - \mathbf{\Sigma}_{FY} \mathbf{\Sigma}_Y{}^{-1} \mathbf{\Sigma}_{FY}^{\mathrm{T}} \tag{3-10}$$

where $\mathbf{\Sigma}_{FY} = \mathbf{\Sigma}_F \mathbf{\Omega}_Y^{\mathrm{T}}$. It should be noted that $\boldsymbol{\mu}_{F|\mathbf{y}}$ is an affine function of observation vector $\mathbf{y}$, while $\mathbf{\Sigma}_{F|Y}$ is a function of the set of measurements Y only; this has important computational implications for sensor placement metrics, as discussed in Chapter 5.

The measurement formulation of Equation 3-5 can describe a variety of measurement topologies. For example, measurements might be local, i.e., of single random field variables. This is accomplished by defining $\mathbf{\Omega}_Y$ such that the row of $\mathbf{\Omega}_Y$ corresponding with the local measurement consists of all zeroes, except for a single one in a position corresponding to the location of the variable to be measured in the vector $\mathbf{f}$. Measurements can also be global, i.e., linear functions of

multiple random field variables across the spatial domain. For example, an average measurement of all random field variables can be obtained by setting each entry in the row of $\mathbf{\Omega}_Y$ corresponding with this measurement to $1/n_X$.

Finally, note that the assumption of Gaussian observations is critical to facilitate closed-form probabilistic inference as in Equation 3-10. In cases where non-Gaussian variables are first transformed to allow for their description as a Gaussian process, observations of these variables, when transformed in the same way, must be Gaussian. Such variables will therefore not be originally Gaussian. For example, in the case of a log-normal random field, observations should be modeled with multiplicative log-normal noise, such that following the logarithm transformation this noise will be additive and Gaussian, as in Equation 3-6.

3.3.3. Affine Limit States

Limit-state variables are assumed to be defined through affine functions of the random field variables:

$$\mathbf{g} = \mathbf{\Xi}_G \mathbf{f} + \mathbf{g}_0 \tag{3-11}$$

where $\mathbf{\Xi}_G$ is the matrix defining which variables or linear combinations of variables in $\mathbf{f}$ contribute to the limit-state variables, and $\mathbf{g}_0$ is a vector of constants allowing for $\mathbf{g}$ to be nonzero even when $\mathbf{f}$ is zero. It is assumed that there are n limit-state variables, one corresponding to each component of the system, and therefore $\mathbf{\Xi}_G$ is a n by n_X matrix.

As an affine function of a vector with a joint Gaussian distribution, the distribution of $\mathbf{g}$ in Equation 2-9 is itself a multivariate Gaussian:

$$\mathbf{g} \sim \mathcal{N}(\boldsymbol{\mu}_G, \mathbf{\Sigma}_G) \tag{3-12}$$

where:

$$\boldsymbol{\mu}_G = \mathbf{\Xi}_G \boldsymbol{\mu}_F + \mathbf{g}_0 \quad \mathbf{\Sigma}_G = \mathbf{\Xi}_G \mathbf{\Sigma}_F \mathbf{\Xi}_G^{\mathrm{T}} \tag{3-13}$$

This definition for the limit-states can be used to describe a number of different functional topologies for the system. For example, a limit-state might be described in terms of a fixed threshold on a single random field variable. In such a case, the corresponding row of $\mathbf{\Xi}_G$ would have only one non-zero entry corresponding to the location of that random variable in $\mathbf{f}$, and the threshold on this variable would be included in $\mathbf{g}_0$. Alternatively, a limit-state might depend on the relative values of several different random fields, for example in the case of an uncertain demand modeled by one random field and an uncertain capacity to resist this demand modeled by another. In this case, positive and negative entries in the row of $\mathbf{\Xi}_G$ would be used to compare these values, and the corresponding entry in $\mathbf{g}_0$ would be zero. Chapter 6 includes several examples of such alternative limit-state definitions under different problem contexts. One example of particular interest in civil engineering is the use of $\mathbf{\Xi}_G$ to encode influence lines,

which transform the loading profile on a structure (vector **f**) to a set of quantities of interest for engineering design, such as moment or shear forces (vector **g**).

3.3.4. Binary Component States

In many cases the management of infrastructure systems is concerned with discrete states, corresponding to certain functional conditions of the component or system in question. The most basic case of this, which will be the primary focus of this work, is the binary state. It is assumed that a component can be in either one of two states. These are generally interpreted as an "active" or "operational" state, denoted by the number one, and a "passive" or "failed" state, denoted by the number zero.

Let the binary state of the ith component be given by:

$$s_i = \mathbb{I}(g_i \geq 0) \tag{3-14}$$

where $\mathbb{I}(\cdot)$ is the indicator function, taking on a value of one when its argument is true and zero otherwise, and s_i and g_i denote the ith entries in **s** and **g** respectively. This definition replaces the probabilistic relationship of Equation 2-10 with a deterministic one. Note that this state is defined based on the corresponding limit-state function being above a threshold, where this threhsold is set arbitrarily to zero; any alternate threshold can be captured through appropriate definition of $\boldsymbol{\Xi}_G$ and $\mathbf{g}_0$ in Equation 3-11.

As each state variable is assumed to be binary, is distribution is Bernoulli:

$$s_i \sim \text{Bernoulli}\left(1 - P_{f,i}\right) \tag{3-15}$$

where $P_{f,i}$ is the probability of failure of the ith component, i.e., $P_{f,i} = \mathbb{P}(g_i < 0)$. The use of the failure probability rather than operational probability to parameterize this distribution is arbitrary, and related to the pessimistic outlook that motivated the use of a loss function rather than a utility function earlier in this work.

Note that, since the probability of failure is descibed using a threshold on the Gaussian random variable g_i, this probability can be expressed as:

$$P_{f,i} = \Phi(-\beta_i) \tag{3-16}$$

where $\Phi(\cdot)$ denotes the standard normal cumulative distribution function, and β_i is the reliability index of the ith component, defined following e.g., Ditlevsen and Madsen (1996) as:

$$\beta_i = \frac{\mu_{g_i}}{\sigma_{g_i}} \tag{3-17}$$

where μ_{g_i} is the mean of g_i, i.e., the ith entry in vector $\boldsymbol{\mu}_G$, and σ_{g_i} is the standard deviation of g_i, i.e., the square root of the ith entry along the diagonal of $\boldsymbol{\Sigma}_G$. The use of the reliability index to relate the distribution of $\mathbf{g}$ to the failure probabilities of the system components will be exploited in Chapter 5.

3.3.5. Binary System States

Finally, as with the component states, the system state (where it is defined as separate from the component states, i.e., in non-cumulative systems) is assumed to be a binary variable, following the same conventions for component states discussed above. The system state is defined as a deterministic function of the component states:

$$\mathbb{s} = \mathbb{S}(\mathbf{s}) \tag{3-18}$$

where the form of the system state function $\mathbb{S}(\cdot)$ is determined by the system's network topology. This replaces the probabilistic relationship of Equation 2-12 with a deterministic one.

Typically in this work, even in non-cumulative systems, the system state will not be explicitly defined. This is a consequence of the fact that, in most problems considered, it is more convenient to define the loss function directly as a function of the component states and managmenet actions, as in Figure 3-5, than to use the system state for this definition. However, in discussing such problems it is often convenient to refer to the system state, even if it is not used computationally. For this reason, two specific forms of the system state function of Equation 3-18 are given in the following subsections for series and parallel systems, the two non-cumulative system network topologies examined in detail later in this work.

3.3.5.1. Series system state

In a series system, the system as a whole will function (i.e., have state value one) if all of its components are functioning. This relationship is captured in the series system's state function:

$$\mathbb{s}_{\text{series}} = \prod_{i=1}^{n} s_i \tag{3-19}$$

3.3.5.2. Parallel system state

In a parallel system, the system as a whole will function if any one of its components is functioning. This relationship is captured in the parallel system's state function:

$$\mathbb{s}_{\text{parallel}} = 1 - \prod_{i=1}^{n}(1 - s_i) \tag{3-20}$$

Chapter 4

The Value of Information Metric

This chapter provides a comprehensive background on the VoI metric, which is the primary sensor placement metric considered in this work. Section 4.1 defines the metric in the context of the infrastructure system management model introduced in Section 2.1, and discusses its use as an objective for guiding optimal sensor placement. In Section 4.2, several prototypical loss functions used for analyzing the VoI in various system topologies considered throughout the remainder of this work are presented. Following from this, Section 4.3 presents a parametric analysis of the VoI provided by a measurement in a system with a single component, illustrating some key properties of the metric in this simplest case of its application.

To examine the relationships between VoI and other common sensor placement metrics, Section 4.4 presents formulations for the prediction error, misclassification rate, and conditional entropy metrics within the system modeling context used in this work. The various properties of these metrics, compared to those of the VoI metric, are examined here, especially in the case of a single-component system. For many of these metrics, an appropriate formulation of the problem loss function can be selected such that the metric can be interpreted as a special case of the VoI.

Finally, Section 4.5 illustrates and discusses the lack of submodularity of the VoI metric in general. Recall from Section 1.2.4 that the property of submodularity allows for theoretical guarantees on the performance of a greedy algorithm for performing combinatorial optimization. Although the VoI is not in general submodular, a greedy approach is still followed in this work for optimizing sensor placement based on this metric. This section discusses specific cases where submodularity does apply, as well as the potential consequences faced when this property does not, and some intuitive safeguards to avoid poor optimization performance.

4.1. Value of Information

The VoI metric is a means for supporting the optimal collection of information in the context of a decision-making problem. This section defines this metric in the context of the infrastructure system modeling framework outlined in Section 2.1. This definition follows the general idea of the VoI metric presented by Raiffa and Schlaifer (1961). For the purposes of this definition, the basic modeling assumptions of Section 3.3 are relaxed in this section.

In managing an infrastructure system, an agent must select the set of actions that he or she will implement from among potential management actions $\mathcal{A}$, as discussed in Section 2.1.4. This decision is guided by the agent's utility or loss function, as presented in Section 2.1.5. A rational decision-maker will choose actions that will minimize the expected value of the loss function across possible states of the random variables affecting the system. The minimum value of the loss that the agent can expect to incur by choosing such actions is designated the prior expected loss:

$$\mathbb{E}\mathrm{L}(\emptyset) = \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_F \mathrm{L}(\mathbf{f}, \mathbf{a}) \tag{4-1}$$

where $\mathbb{E}_F$ denotes the statistical expectation with respect to the prior probability distribution of the random field variables p_F. That is, without additional information, the decision-maker will select a set of actions that will minimize the expected value of the loss function across the prior distribution of possible values of the random field variables.

Note that for alternative formulations of the loss as a function of $\mathbf{g}$, $\mathbf{s}$, or $\mathbb{s}$, this expectation can be calculated using the appropriate conditional distributions introducted in Section 2.1. For example, for a loss function forumated as $\mathrm{L}(\mathbf{g}, \mathbf{a})$, the relationship of Equation 2-15 can be substituted into Equation 4-1. For the sake of generality, the definitions presented in this section assume a loss function of the form $\mathrm{L}(\mathbf{f}, \mathbf{a})$.

The prior expected loss is associated with a set of prior optimal actions that will, in an expected sense, lead to this minimum loss value for the decision-making agent. The set of prior optimal actions is denoted:

$$\mathbf{a}^*(\emptyset) = \mathrm{argmin}_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_F \mathrm{L}(\mathbf{f}, \mathbf{a}) \tag{4-2}$$

If, instead of selecting decisions in the absence of additional information (other than the prior random field model for the system), the decision-making agent has access to observations $\mathbf{y}$ obtained for measurement set Y, he or she should instead choose their actions in order to minimize the expected value of the loss under the posterior distribution for the random field variables. This value is denoted as the conditional expected loss, i.e., the expected loss conditional to an observation $\mathbf{y}$ of Y:

$$\mathbb{E}\mathrm{L}(\mathbf{y}) = \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}} \mathrm{L}(\mathbf{f}, \mathbf{a}) \tag{4-3}$$

The posterior optimal action choice that leads (in an expected sense) to the conditional expected loss is:

$$\mathbf{a}^*(\mathbf{y}) = \mathrm{argmin}_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}} \mathrm{L}(\mathbf{f}, \mathbf{a}) \tag{4-4}$$

While $\mathbb{E}\mathrm{L}(\mathbf{y})$ denotes the expected loss conditional to observation $\mathbf{y}$, this observation is itself a random variable. Accounting for this necessitates pre-posterior analysis, i.e., the expectation of

the conditional expected loss must be taken across possible outcomes of the measurement activities. The result is the posterior expected loss:

$$\mathbb{E}\mathrm{L}(Y) = \mathbb{E}_Y \mathbb{E}\mathrm{L}(\mathbf{y}) = \mathbb{E}_Y \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}} \mathrm{L}(\mathbf{f}, \mathbf{a}) \tag{4-5}$$

Intuitively, the posterior expected loss is the expected value of the loss function incurred by the agent managing the system, considering that he or she will have access to measurements Y of the random field variables F to help guide the choice of optimal management actions for the system.

The VoI for measurement set Y is defined as the difference of the prior and posterior expected losses:

$$\mathrm{VoI}(Y) = \mathbb{E}\mathrm{L}(\emptyset) - \mathbb{E}\mathrm{L}(Y) \tag{4-6}$$

This expresses the expected savings, in terms of the reduction in the expected loss, resulting from an observation of Y being available before the management actions are decided.

4.1.1. Bounds on the Value of Information

The VoI has several important properties that define its boundaries. First, it is strictly non-negative. This is a consequence of the "information never hurts" principle; for a rational decision-maker, incorporating additional information into the decision-making process can only reduce, in an expected sense, the management costs for the system (Hilton, 1977). Therefore, zero is the natural lower-bound for $\mathrm{VoI}(Y)$. Furthermore, VoI is upper-bounded by the value of complete perfect information (VoPI), which is evaluated as $\mathrm{VoI}(F)$. That is, VoPI is achieved when the system behavior is known exactly through an observation of all random field variables F prior to taking actions, i.e., when the system is no longer uncertain. The bounds on $\mathrm{VoI}(Y)$ are therefore:

$$0 \leq \mathrm{VoI}(Y) \leq \mathrm{VoI}(F) \tag{4-7}$$

Appendix A includes derivations of these bounds for the interested reader.

4.1.2. Regret-based Formulation

An alternative but equivalent method for formulating the VoI that is of interest in some applications is based on a regret function, rather than a loss function. The regret function measures, for any specific system management outcome, how much better the best possible outcome in that situation would be. Formally, regret can be defined from the loss function as:

$$\mathrm{R}(\mathbf{f}, \mathbf{a}) = \mathrm{L}(\mathbf{f}, \mathbf{a}) - \mathrm{L}\big(\mathbf{f}, \mathbf{a}^*(\mathbf{f})\big) \tag{4-8}$$

where:

$$\mathbf{a}^*(\mathbf{f}) = \mathrm{argmin}_{\mathbf{a}\in\mathcal{A}}\, \mathrm{L}(\mathbf{f},\mathbf{a}) \tag{4-9}$$

That is, the regret is the difference between the loss and the minimum possible loss suffered when choosing the posterior optimal action given knowledge of the random field variables. A useful property of this formulation is that the regret is lower-bounded by zero, and takes this value if and only if the posterior optimal action is chosen.

Under the regret formulation, the prior and posterior expected regret are defined similar to Equations 4-1 and 4-5 as:

$$\mathbb{E}\mathrm{R}(\emptyset) = \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_F \mathrm{R}(\mathbf{f},\mathbf{a}) \quad \mathbb{E}\mathrm{R}(Y) = \mathbb{E}_Y \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}} \mathrm{R}(\mathbf{f},\mathbf{a}) \tag{4-10}$$

The VoI is defined similar to Equation 4-6 as:

$$\mathrm{VoI}(Y) = \mathbb{E}\mathrm{R}(\emptyset) - \mathbb{E}\mathrm{R}(Y) \tag{4-11}$$

These two formulations of the VoI are equivalent; a proof is provided in Appendix A. The loss-based formulation will be adhered to in this work, except where otherwise indicated.

4.1.3. Sensor Placement using Value of Information

VoI corresponds to the maximum amount one should be willing to pay for obtaining information before making a decision. That is, if $\mathrm{C}(Y)$ is the cost of measuring Y, then $\mathbb{E}\mathrm{L}(Y) + \mathrm{C}(Y)$ should be less than $\mathbb{E}\mathrm{L}(\emptyset)$ for the measurements to be cost-effective. This observation motivates the definition of an objective for sensor placement based on the VoI:

$$\mathrm{M}_{\mathrm{VoI}}(Y) = \mathrm{VoI}(Y) - \mathrm{C}(Y) \tag{4-12}$$

This objective measures the quality of sensor placement set Y as the VoI resulting from the measurement of Y minus the cost (measured in the same scale as the loss function, e.g., a monetary scale) of obtaining this measurement. This is also designated the net VoI of Y.

The objective of Equation 4-12 can be used directly in the optimization problem of Equation 2-16. It can also be used directly for greedy optimization of sensor placements in the offline case, as in Equation 2-17. For online sensor placement, the conditional objective of Equation 2-18 using VoI is:

$$\mathrm{M}_{\mathrm{VoI}}\left(Y_j|\mathbf{y}_{j-1}\right) = \mathrm{VoI}\left(Y_j|\mathbf{y}_{j-1}\right) - \mathrm{C}\left(\left\{Y_1^*,\dots,Y_{j-1}^*,Y_j\right\}\right) + \mathrm{C}\left(\left\{Y_1^*,\dots,Y_{j-1}^*\right\}\right) \tag{4-13}$$

where:

$$\mathrm{VoI}\left(Y_j|\mathbf{y}_{j-1}\right) = \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}_{j-1}} \mathrm{L}(\mathbf{f},\mathbf{a}) - \mathbb{E}_{Y_j|\mathbf{y}_{j-1}} \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\{\mathbf{y}_{j-1},y_j\}} \mathrm{L}(\mathbf{f},\mathbf{a}) \tag{4-14}$$

This formulation describes the incremental VoI, i.e., the additional value provided by measurement Y_j when measurements $Y_1^*,\dots,Y_{j-1}^*$ have already been observed. This incremental

VoI should be greater than the incremental cost of adding the observation to the set for it to be selected by the algorithm.

Due to the need to only take the expectation over Y_j rather than over $\{Y_1^*, \dots, Y_{j-1}^*, Y_j\}$ for evaluating the outer expectation for the posterior expected loss in computing VoI, online greedy optimization is computationally simpler than offline placement. Furthermore, for the same system, the objective gain (in terms of reduction in the posterior expected loss) will be higher in the online case than in the offline, due to the inclusion of actual collected data about the underlying random field when performing the optimization (see Appendix A for details). However, note that the greedy algorithm presented here presents a myopic approach to online placement, since each selected placement is assumed to be the final placement, and VoI is evaluated accordingly, i.e., without considering that this information will also be used to help guide the future steps of the online placement algorithm. Such a myopic approach can be suboptimal in general (Heckerman et al., 1993). Finally, note that $\mathbb{E}\mathrm{L}(\emptyset)$ is constant with respect to Y, and so in evaluating Equation 4-6 when performing offline sensor placement optimization, this term does not impact the optimization objective, and therefore can be ignored to simplify computation.

4.2. Prototypical Loss Functions

This section introduces the prototypical loss function formulation for infrastructure management that is used extensively in this work. This function is based on the assumption of binary component states, as in Section 3.3.4, and binary management actions for each component. These actions are broadly described as a "repair" action, i.e., an active intervention to correct a potential failure of a component, and a "do nothing" action. These actions are denoted by the selection of $a_i = 1$ for the "repair" action and $a_i = 0$ for the "do nothing" action for the ith component. The "repair" action is associated with a fixed repair cost, denoted C_r, which must be paid by the agent to undertake this action. The "do nothing" action is assumed to be free; however, the agent must face the potential consequences of suffering a failure of the component or system due to the lack of intervention. This consequence is summarized by a failure cost C_f, which accounts both for the direct consequences in terms of system damage and down-time, as well as the indirect costs to society of failing to have the component or system operating as expected due to this failure.

This prototypical management problem represents the simplest decision-making problem of an agent: whether or not to take a possible active intervention action, and thus pay a known penalty (the repair cost), or take no action, and have the consequence determined by the uncertain behavior of the system. The appropriate choice will depend on the relative costs of the outcomes, the probability of a component or system failure, and (in the posterior case) the availability of appropriate information collected within the system to help support decision-making. Most decision-making problems examined in this work will be adapted to this prototypical form.

The following sub-sections formulate loss functions based on this prototypical decision-making problem in the context of different system topologies, including cumulative and non-cumulative (especially series and parallel) systems.

4.2.1. Cumulative Systems

In a cumulative system, the prototypical decision-making problem is posed separately for each system component. The component-level loss function within the decomposable loss function of Equation 3-1 is expressed in the following way for the ith component:

$$\mathrm{L}_i(s_i, a_i) = C_{f,i}(1 - s_i)(1 - a_i) + C_{r,i}a_i \tag{4-15}$$

where $C_{f,i}$ is the cost of the component's failure, $C_{r,i}$ is the cost to repair or replace the component, and s_i and a_i are the binary state and action for managing the component, respectively. As a direct consequence of the decomposability of the loss function assumed for the cumulative system topology, the cost for managing the system is the sum of these component-level costs, i.e., the consequences of a component failure or the choice to repair or replace a component have no effect on the performance of the other system components.

4.2.2. Non-Cumulative Systems

In general, for systems without a cumulative topology, consequences of component failures may extend to the system as a whole, e.g., through impacts on the system network structure. Furthermore, repair costs may not be additive, e.g., the cost of repair per component may be reduced when multiple components are repaired due to economies of scale or the fixed cost to mobilize repair teams regardless of the number of repairs.

General loss function forms are presented for two specific non-cumulative system cases. First, in the case that a single management action applies to the entire system, the loss function is formulated as:

$$\mathrm{L}(\mathbf{s}, a) = C_f\big(1 - \mathbb{S}(\mathbf{s})\big)(1 - a) + C_r a \tag{4-16}$$

where C_f is the cost of system failure and C_r is the cost of performing the system-level repair action, i.e., of selecting $a = 1$. Note that, for consistency with the formulation of Figure 3-5, the system state function of Equation 3-18 is incorporated directly into the loss function. This loss function captures situations in which a common action must be chosen for the entire system, rather than for each component as in the cumulative system case above. For example, evacuation, closure, or quarantine orders are issued for entire regions or systems, rather than single components or individuals.

Second, it is considered that actions are taken for components separately, as in the cumulative system topology, but that the consequences of component failures are assessed at the system-level. Under this assumption, the loss function is:

$$\mathrm{L}(\mathbf{s},\mathbf{a}) = C_f\big(1 - \mathbb{S}'(\mathbf{s},\mathbf{a})\big) + \sum_{i=1}^{n} C_{r,i} a_i \tag{4-17}$$

where C_f is the cost of system failure and $\mathbb{S}'(\mathbf{s},\mathbf{a})$ is the effective state of the system, taking into account the possibility that component-level repair actions will reverse a potential system-level failure. Two formulations for this effective state in series and parallel systems are shown below. Finally, note that it is assumed in this case that the total cost for undertaking management actions is the sum of these costs for each component in the system (expressed by the summation term in the second part of Equation 4-17). In general, this assumption may not hold, and the repair cost may also be a nonlinear function of the component-level actions.

4.2.2.1. Series systems

In a series system, motivated by the state function of Equation 3-19, the effective system state taking repair actions into account is:

$$\mathbb{S}'_{\text{series}}(\mathbf{s},\mathbf{a}) = \prod_{i=1}^{n}\big(1 - (1 - s_i)(1 - a_i)\big) \tag{4-18}$$

That is, the system will function if either all components are functioning, or if, for everey non-functional component, a repair action is selected.

4.2.2.2. Parallel systems

In a parallel system, motivated by the state function of Equation 3-20, the effective system state taking repair actions into account is:

$$\mathbb{S}'_{\text{parallel}}(\mathbf{s},\mathbf{a}) = 1 - \prod_{i=1}^{n}(1 - s_i)\,(1 - a_i) \tag{4-19}$$

That is, the system will function if at least one component is functioning, or if at least one component has been repaired or replaced.

4.3. Parametric Analysis on a Single-Component System

Factors such as sensor precision, management costs, and prior component reliability can affect the VoI obtained from a measurement on a component in complicated ways. To begin to investigate these effects, this section presents a parametric analysis of the VoI in a single-component system, following the prototypical decision-making problem formulation of Section 4.2. Note that, for a system with only one component, the topology distinctions mentioned in that section do not apply.

In this analysis, the state of a single component is affected by a single Gaussian random variable f, modeled as in Equation 3-4 with mean μ_f and standard deviation σ_f. One potential observation is considered for this component: a direct observation of f with zero-mean Gaussian noise, i.e., y is defined as in Equation 3-5 with $\mathbf{\Omega}_Y = 1$, and error $\boldsymbol{\epsilon}$ is univariate with mean 0 and standard deviation σ_ϵ. The limit-state variable is defined directly as $g = f$, i.e., $\mathbf{\Xi}_G = 1$ and $\mathbf{g}_0 = 0$ in Equation 3-11. The component's state s is defined as in Equation 3-14.

A binary decision-making problem is defined for this component, following the prototypical problem formulation of Section 4.2. The loss function for this component is given, following the form of Equation 4-15, as:

$$L(s, a) = C_f(1 - s)(1 - a) + C_r a \tag{4-20}$$

This prototypical component management problem involves the choice between doing nothing, with the potential to incur loss C_f if the component fails, and repairing the component at cost C_r.

To illustrate this decision-making problem, Figure 4-1a plots the expected value of the loss function for the component versus the probability of failure, computed for the component as in Equation 3-16, under the two possible management actions for this component. For the purposes of this figure, a failure cost of $C_f = \$300\text{M}$ and a repair cost of $C_r = \$50\text{M}$ are assumed. Note that the expected loss under the "do nothing" action increases linearly as the probability of failure increases:

$$\mathbb{E}_S \mathrm{L}(s, a = 0) = C_f \mathbb{P}(s = 0) = C_f P_f \tag{4-21}$$

The expected loss under the repair action is constant at C_r.

With a known prior probability of failure, the prior optimal action would be the choice that minimizes the expected loss. The expected loss taking the optimal action is indicated by the solid line in Figure 4-1a. Note here that if $P_f < 16.67\%$, the prior optimal choice is to do nothing, while if $P_f > 16.67\%$, the prior optimal choice is to repair the component. The value for P_f at which the choice of prior optimal action changes, i.e., the value at which the prior expected loss of both decisions is equal, is designated the *decision point*, and is evaluated as:

$$P^* = \frac{C_r}{C_f} \tag{4-22}$$

It is assumed here that $P^* \leq 1$; if this were not the case, the choice to do nothing would always be optimal, and there would be no need for any further decision-making analysis. Using this notation, the prior expected loss for the component, following Equation 4-1 with the loss function of Equation 4-20, can be expressed as:

$$\mathbb{E}\mathrm{L}(\emptyset) = \min_{\mathbf{a} \in \{0,1\}} \left[C_f P_f (1 - a) + C_f P^* a \right] = C_f \min\left\{ P_f, P^* \right\} \tag{4-23}$$

For the example of Figure 4-1, the component's prior probability of failure is assumed to be $P_f = 10\%$, and so the prior optimal action is to do nothing, with a prior expected loss of \$30M.

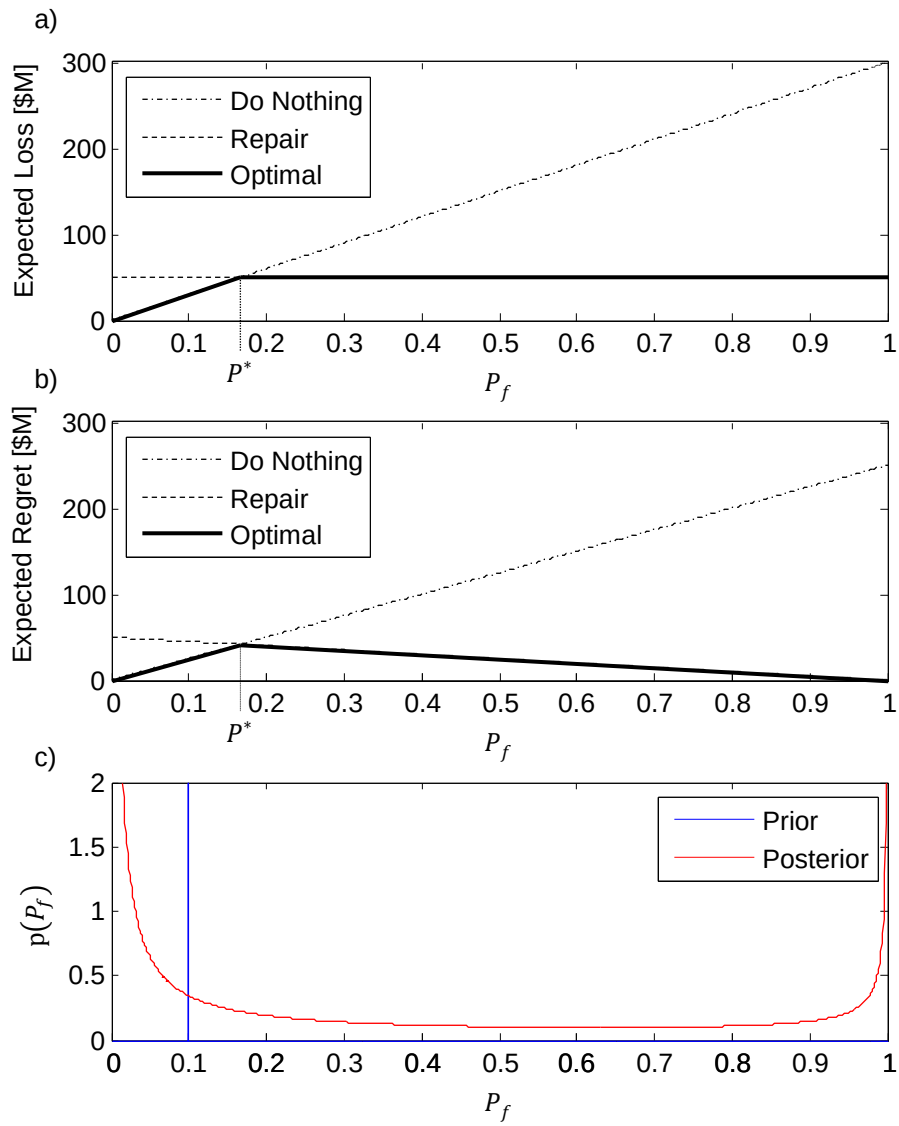

Figure 4-1: a) Example expected loss versus failure probability for management of a single component; b) Example expected regret versus failure probability for management of a single component; c) Example prior and posterior failure probability distributions for a single component.

Figure 4-1b illustrates the same decision-making problem as Figure 4-1a using the regret instead of the loss. Note that the expected regret takes on value 0 at either extreme for P_f, i.e., if the behavior of the component is known with either $P_f = 0\%$ or $P_f = 100\%$, the optimal action can

be selected with zero regret. The value of P^* is unchanged by shifting between loss and regret formulations.

Following the Gaussian updating procedure from Equation 3-10, the posterior distribution parameters for f conditional to observation y are:

$$\mu_{f|y} = \mu_f + \frac{\sigma_f^2}{\sigma_f^2+\sigma_\epsilon^2}\left(y - \mu_f\right) = \frac{r^2\mu_f+y}{1+r^2} \quad \sigma_{f|y} = \sqrt{\sigma_f^2 - \frac{\sigma_f^4}{\sigma_f^2+\sigma_\epsilon^2}} = \frac{r\sigma_f}{\sqrt{1+r^2}} \tag{4-24}$$

where $r = \sigma_\epsilon/\sigma_f$ is the measurement noise ratio, i.e., the ratio of the standard deviation of the additive noise to the standard deviation of the value being measured. Note that the observation y is itself a Gaussian random variable with mean $\mu_y = \mu_f$ and standard deviation $\sigma_y = \sqrt{\sigma_f^2 + \sigma_\epsilon^2} = \sigma_f\sqrt{1+r^2}$. Conditional to this observation, the posterior probability of failure for the component is computed, following Equations 3-16 and 3-17, as:

$$P_{f|y} = \Phi\left(-\frac{\mu_{f|y}}{\sigma_{f|y}}\right) \tag{4-25}$$

The conditional expected loss of Equation 4-3, following the form of Equation 4-23, can then be expressed as:

$$\mathbb{E}\mathrm{L}(y) = \min_{\mathbf{a}\in\{0,1\}}\left[C_f P_{f|y}(1-a) + C_f P^* a\right] = C_f \min\left\{P_{f|y}, P^*\right\} \tag{4-26}$$

The VoI for this measurement is therefore assessed, following Equations 4-5 and 4-6, as:

$$\mathrm{VoI}(Y) = C_f\left[\min\left\{P_f, P^*\right\} - \mathbb{E}_Y \min\left\{P_{f|y}, P^*\right\}\right] \tag{4-27}$$

Note from Equation 4-24 that $\sigma_{f|y}$ is a constant, and that $\mu_{f|y}$ is an affine function of observation y. Since y has a Gaussian distribution, the argument of Equation 4-25 will have a Gaussian distribution, and therefore the posterior component failure probability $P_{f|y}$ is itself a random variable. Due to the nonlinear nature of the Gaussian cumulative distribution function $\Phi(\cdot)$, the distribution of $P_{f|y}$ will not be Gaussian.

Figure 4-1c illustrates possible prior and posterior distributions for the probability of component failure for this single component example. The prior distribution in this figure reflects the example prior failure probability of $P_f = 10\%$ with a Dirac delta distribution at this value. A posterior distribution is shown for a measurement y with $r = 0.25$. This posterior distribution is obtained by processing the posterior distribution for f conditional to y through the non-linear transformation of Equation 4-25. The result is a bimodal posterior distribution for $P_{f|y}$, as shown in the figure. Note that, in the extreme case of $r = 0$ (i.e., where y is a direct observation of f), this distribution will consist of two point probability masses: a mass of 10% at $P_{f|y} = 1$ and a

mass of 90% at $P_{f|y} = 0$, i.e., there is a 10% chance that a failure will occur, and conditioned on a measurement of this failure the posterior probability of failure is 100%, while there is a 90% chance the component will be operational, and conditional to this measurement the posterior probability of failure is 0.

The fact that VoI is non-negative is clear from Figure 4-1 as a consequence of Jensen's Inequality. For both the loss of Figure 4-1a and the regret of Figure 4-1b, these functions are concave with respect to the failure probability. Jensen's Inequality for concave functions states that the expected value of the function over its argument is less than or equal to the function evaluated at the argument's expected value (Jensen, 1906). In this case, the argument is the failure probability, and its expected value is the prior failure probability P_f. The function evaluated at its expected argument value is therefore the prior expected loss $\mathbb{E}\mathrm{L}(\emptyset)$. The expectation of the function over its argument is the expected value of the conditional loss $\mathbb{E}\mathrm{L}(y)$ over the distribution of the posterior failure probability $P_{f|y}$, or $\mathbb{E}\mathrm{L}(Y)$. By Jensen's Inequality, the latter quantity is less than or equal to the former, i.e., $\mathbb{E}\mathrm{L}(Y) \leq \mathbb{E}\mathrm{L}(\emptyset)$, proving the non-negativity of the VoI in this case.

From Equation 4-27, it can be determined that VoI in a single-component system is a function of four parameters:

1. The cost of component failure C_f
2. The decision point P^* (as determined from C_f and C_r via Equation 4-22)
3. The prior component probability of failure P_f (as determined from μ_f and σ_f via Equations 3-16 and 3-17)
4. The measurement noise ratio r (which, together with μ_f and σ_f, defines the posterior distribution for $P_{f|y}$ via Equations 4-24 and 4-25)

From the form of Equation 4-27, it is clear that VoI is linear with respect to C_f. However, the relationship between VoI and the other parameters is less obvious. For this reason, a parametric analysis is performed, varying P^*, P_f, and r and evaluating the resulting VoI for the one-component system. Figure 4-2 demonstrates these results via two cross-sections of the three-dimensional parametric analysis space. In the figure, to present a more reasonable scale, P_f is represented by the reliability index β; these are related (following Equation 3-16) as $\beta = -\Phi^{-1}(P_f)$. Furthermore, P^* is represented by β^*; these are related as $\beta^* = -\Phi^{-1}(P^*)$.

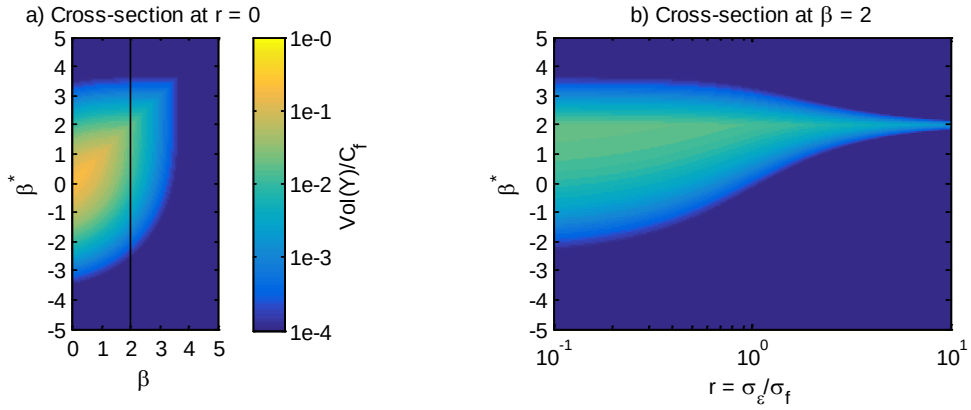

Figure 4-2: Parametric analysis results for VoI in a single-component system, with cross-sections for a) $r = 0$; and b) $\beta = 2$.

In Figure 4-2a, the scaled VoI, $\mathrm{VoI}(Y)/C_f$, is plotted versus β and β^* for a measurement noise level of $r = 0$, i.e., under the assumption of a perfect observation where $y = f$. This represents the VoPI for this problem, which can be expressed in closed-form as:

$$\mathrm{VoPI} = C_f\left[\min\{P_f, P^*\} - P_f P^*\right] \tag{4-28}$$

where the expectation over Y in Equation 4-27 is handled by noting that, as discussed above, $P_{f|y} = 0$ with probability $1 - P_f$ and $P_{f|y} = 1$ with probability P_f for a perfect observation. Note that Figure 4-2a is symmetric about the point $\beta = 0$, $\beta^* = 0$.

In Figure 4-2b, the scaled VoI is plotted versus β^* and r for a fixed $\beta = 2$, corresponding to a prior failure probability of about 2%. Note that, regardless of the noise level r, VoI is maximized when $\beta^* = \beta$, i.e., when the prior failure probability is at the decision point. This corresponds to the point in Figure 4-1b where prior expected regret is maximized. Therefore, in this case, additional information has the greatest opportunity to reduce the regret, and therefore the VoI is greatest. Also note that measurements with a noise ratio greater than about 1 provide relatively little VoI unless β is very close to β^*.

Finally, consider that observation y is not a direct measure of f, but rather a measure of another random field variable that is correlated with f. That is, consider that the component in question is located at $\mathbf{x}_1$, and is affected by $f_1 = \mathrm{f}(\mathbf{x}_1)$, and that the measurement y of the random field is obtained at $\mathbf{x}_2$, where $\mathrm{k}(\mathbf{x}_1, \mathbf{x}_2) = \sigma_{f_1}\sigma_{f_2}\rho_{12}$, with noise level $r_2 = \sigma_\epsilon/\sigma_{f_2}$. Following Equation 3-10, the posterior standard deviation for f_1 given y is now:

$$\sigma_{f_1|y} = \sqrt{\sigma_{f_1}^2 - \frac{\rho_{12}^2 \sigma_{f_1}^2 \sigma_{f_2}^2}{\sigma_{f_2}^2 + \sigma_\epsilon^2}} = \sigma_{f_1} \sqrt{\frac{1 + r_2^2 - \rho_{12}^2}{1 + r_2^2}} \tag{4-29}$$

From this relationship, an equivalent noise level $\tilde{r}$ is computed such that the reduction in the posterior standard deviation of f_1 due to this indirect observation is the same as if an observation with noise level $\tilde{r}$ were made on f_1 directly. Substituting $\tilde{r}$ for r in Equation 4-24, setting this equal to Equation 4-29, and solving for $\tilde{r}$ gives:

$$\tilde{r} = \sqrt{\frac{1 + r_2^2 - \rho_{12}^2}{\rho_{12}^2}} \tag{4-30}$$

Note that if the observed random field variable is uncorrelated with f_1, i.e., if $\rho_{12} = 0$, $\tilde{r}$ is infinite, corresponding to an uninformative measurement. On the other hand, if f_1 and f_2 are perfectly correlated, i.e., if $\rho_{12} = 1$ (or perfectly negatively correlated, with $\rho_{12} = -1$), then $\tilde{r} = r_2$, i.e., the noise level is as if the measure were made on f_1 directly. Using this method, indirect observations of the random field variable affecting a component can be treated as equivalent direct observations with an appropriately transformed precision. The VoI provided by indirect measures can then be evaluated by substituting the equivalent precision $\tilde{r}$ for r in Figure 4-2.

In summary, this section has investigated the performance of the VoI metric in a system consisting of one component, under the basic modeling assumptions stated in Section 3.3 and under a prototypical binary decision-making problem as described in Section 4.2. In this case, the VoI for a single measurement of the underlying random field can be parameterized by four quantities: the cost of component failure, the ratio of repair and failure costs, the prior probability of component failure, and the relative precision of the measurement. This illustrates how the VoI metric accounts for the economic considerations of the decision-making problem, the prior model for the performance of the component, and the quality of the information used in supporting decision-making. One important result of this section which will be re-visited in Chapter 5 is that a posterior distribution over the probability of failure for a component can be defined for binary components, which allows for computation of the posterior expected loss via an expectation over the univariate posterior failure probability for that component, rather than over a possibly multivariate distribution for Y if multiple observations are considered. This result is used as the basis for an efficient method for evaluating VoI in cumulative systems in Section 5.2.2.

4.4. Relationship to Other Sensor Placement Metrics

The VoI metric for sensor placement is one of many alternative metrics, as discussed in Section 1.2.4. However, it shares common properties with several of these metrics, and certain metrics can be understood as special cases of the VoI. In this section, the prediction error,

misclassification rate, and conditional entropy metrics for sensor placement are introduced in the context of infrastructure management and compared to the VoI metric. The basic modeling assumptions of Section 3.3 are relaxed in this section in order to provide general definitions of the metrics.

4.4.1. Prediction Error

Many basic problems in sensing and machine learning can be characterized as problems of predicting or estimating an unknown quantity using other known or measurable quantities. In infrastructure management, such problems have applications in control scenarios, where prediction of continuous quantities is important to determining appropriate control inputs to restore a system to a desired operating condition (e.g., Anderson and Moore, 2007).

For optimal prediction of continuous quantities, an intuitive approach is to measure quantities that most reduce the expected prediction error (e.g., Krause et al., 2008b). Prediction error can be measured as the absolute difference between the prediction and the true value. However, it is often more convenient to use the squared error, or the ℓ^2 norm of the error, as this quantity has many useful properties, e.g., allowing for computationally efficient solutions for linear regression problems (Berger, 1993). Thus, a potential objective for optimizing sensor placement is to select measures that will reduce the expected squared prediction error.

Such a problem can be expressed as a special case of the VoI metric within the infrastructure system management context. Let the random field variables F be the quantities that are to be predicted, and let the actions A denote the predictions for these quantities. The squared prediction error can thus be assessed using a loss function with a quadratic form:

$$\mathrm{L}_{\mathrm{pred}}(\mathbf{f},\mathbf{a}) = (\mathbf{f}-\mathbf{a})^{\mathrm{T}}(\mathbf{f}-\mathbf{a}) \tag{4-31}$$

This loss function quantifies the quality of the prediciton, and can be used directly in the VoI formulation of Section 4.1 to provide an objective for sensor placement for minimizing the square prediction error.

Note that this loss function also happens to be in a decomposable form, as in Equation 3-1:

$$\mathrm{L}_{\mathrm{pred}}(\mathbf{f},\mathbf{a}) = \sum_{i=1}^{n}(f_i - a_i)^2 \tag{4-32}$$

where the number of components n is the number of random field variables n_X being predicted. This loss function therefore defines a cumulative system topology for random field prediction.

Prediction for transformations of the underlying random field can be addressed in a similar manner, using a definition for the loss function based on an approrpatiely defined limit-state variable **g**:

$$\mathrm{L}_{\mathrm{pred}}(\mathbf{g},\mathbf{a}) = (\mathbf{g}-\mathbf{a})^{\mathrm{T}}(\mathbf{g}-\mathbf{a}) \tag{4-33}$$

This formulation is useful in cases where prediction is of interest not for values of the random field directly, but for transformations of the field, e.g., regional average values or differences between random field quantities.

Finally, note that in this case the loss function is not measured directly in monetary terms, as in the traditional formulation of VoI in Section 4.1. Instead, it is measured directy in terms of the squared unit values of the quantities to be predicted. Therefore, when applying the objective of Equation 4-12, it may be inapprorpate to directly subtract the sensing cost from the VoI. Rather, an equivalence might be established between cost and the prediction error, or cost can be removed entirely from Equation 4-12. Another alternative is to use a benefit-cost ratio formulation:

$$\mathrm{M}_{\mathrm{B/C}}(Y) = \frac{\mathrm{VoI}(Y)}{\mathrm{C}(Y)} \tag{4-34}$$

Such a formulation is not recommended when VoI is expressed in monetary terms directly, since it is much more appropriate in that case to subtract costs from benefits. However, when these quantities are expressed in incompatible units, as in the case of prediction error, this formulation presents a reasonable alternative for trading off benefit and cost in a sensing optimization problem. The use of such a benefit-cost ratio is motivated by the work of Krause (2008), where a similar formulation is used for greedy optimization of the conditional entropy and mutual information metrics, which are also assessed in non-monetary terms. The use of such a ratio can avoid a potential sub-optimal approach via greedy optimization of selecting a single high-cost measurement which provides moderate benefits instead of selecting multiple low-cost measures which each provide little benefit, but when combined can provide a high benefit. The use of such a ratio therefore represents both an objective for sensing, in terms of maximizing the benefit-cost ratio of the overall measurement set, and an algorithmic tool to avoid greedy selection of a small number of high-cost measures which provide lower overall benefit than a larger set of lower-cost measures.

4.4.1.1. Weighted prediction error

In addition to direct prediction of the underlying random field variables, weighted prediction can be considered, where prediction errors are penalized differently for different variables. A reason for adopting this approach might be that certain variables, or only a subset of the underlying random variables, are of interest. Alternatively, this weighting can be used to establish an equivalence between predictions of different quantities measured in different units, or between these predictions and a monetary cost, such that VoI resulting from the prediction-based loss function can be traded off against cost directly as in Equation 4-12.

Weighted prediction error can be handled by the following loss function used to define the decision-making problem:

$$\mathrm{L}_{\mathrm{pred}}(\mathbf{f},\mathbf{a}) = (\mathbf{f}-\mathbf{a})^{\mathrm{T}}\mathbf{Z}^{\mathrm{T}}\mathbf{Z}(\mathbf{f}-\mathbf{a}) \tag{4-35}$$

where $\mathbf{Z}$ is a weight matrix defined appropriately to capture the desired prediction penalties. This matrix is multiplied by its transpose to ensure symmetry of the resulting matrix.

Note that, under the linear relationships between F and G of Section 3.3, this formulation is related to that of Equation 4-33, assuming $\mathbf{g}_0 = \mathbf{0}$ and $\mathbf{Z} = \boldsymbol{\Xi}_G$, as:

$$(\mathbf{g}-\mathbf{a})^{\mathrm{T}}(\mathbf{g}-\mathbf{a}) = (\mathbf{f}-\mathbf{a})^{\mathrm{T}}\boldsymbol{\Xi}_G^{\mathrm{T}}\boldsymbol{\Xi}_G(\mathbf{f}-\mathbf{a}) + 2\mathbf{f}^{\mathrm{T}}\boldsymbol{\Xi}_G^{\mathrm{T}}\boldsymbol{\Xi}_G\mathbf{a} - 2\mathbf{f}^{\mathrm{T}}\boldsymbol{\Xi}_G^{\mathrm{T}}\mathbf{a} + \mathbf{a}^{\mathrm{T}}\mathbf{a} - \mathbf{a}^{\mathrm{T}}\boldsymbol{\Xi}_G^{\mathrm{T}}\boldsymbol{\Xi}_G\mathbf{a} \tag{4-36}$$

In Section 5.2.1, it is shown that sensor placement based on weighted prediction error using this definition of the weight matrix is the same as that for prediction of the limit-state variables. The weighted prediction error formulation is used in the context of temperature prediction in Chapter 10, where the vulnerability of populations to extreme heat is used to define the weighting.

4.4.2. Misclassification Rate

When prediction is performed for discrete rather than continuous quantities, the problem is commonly described as one of classification. A natural metric to use in this context is the misclassification rate, i.e., the probability that the prediction of the discrete state will not be correct (e.g., Rasmussen and Williams, 2006).

This metric can also be formulated as a special case of the VoI, with an appropriate misclassification-based loss. If component states S are the variables that are to be correctly predicted, this loss function is:

$$\mathrm{L}_{\mathrm{MCR}}(\mathbf{s},\mathbf{a}) = \sum_{i=1}^{n} \mathbb{I}(s_i \neq a_i) \tag{4-37}$$

Note that this is an inherently decomposable loss function, as in Equation 3-1, and therefore defines a cumulative system topology for the classification problem. Also note that this loss is formulated here as the total misclassification rate for all components. It can be expressed as the average misclassification rate among components by dividing the loss by n; however, as this is only scaling the loss function by a constant, it will have no impact on the optimization problem.

If prediction of the system state as a whole is of direct interest, rather than prediction of component states individually, the misclassification loss can be formulated as a function of the system state instead:

$$\mathrm{L}_{\mathrm{MCR}}(\mathbb{s},a) = \mathbb{I}(\mathbb{s} \neq a) \tag{4-38}$$

As with the prediction error, the misclassification rate may not trade off directly with cost, and so equivalence between misclassification and cost might be established, or a benefit-cost formulation of the objective as in Equation 4-34 may be adopted.

Finally, referring back to the binary component of Section 4.3 and the expected loss as a function of component failure probability as shown in Figure 4-1a, the misclassification rate can also be plotted as a function of failure probability, as shown in Figure 4-3. Note that the misclassification rate takes on its peak value of 0.5 at $P_f = 50\%$, where the binary state of the component is the most uncertain. Also note the similarity to the expected regret of Figure 4-1b; misclassification loss is a special case of the regret in this problem when the cost of a failure is set to twice the cost of a repair.

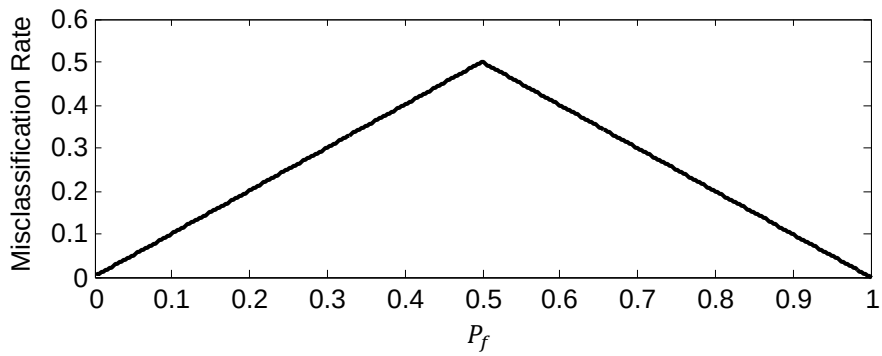

Figure 4-3: Misclassification rate as a function of failure probability for a binary component.

4.4.3. Conditional Entropy

Entropy is an information theoretic measure of the uncertainty of a set of random variables (Cover and Thomas, 2006). Intuitively, the goal of sensing might be understood as the reduction in uncertainty in one set of variables, conditional to measurements of these or other related variables (e.g., Currin et al., 1991). In the context of infrastructure systems, the entropy of the random field variables affecting the system is denoted $\mathrm{H}(F)$, and is computed as:

$$\mathrm{H}(F) = -\mathbb{E}_F \log\left(\mathrm{p}_F(\mathbf{f})\right) \tag{4-39}$$

This represents the marginal or prior entropy of F, as it is derived from the prior distribution p_F. The posterior entropy conditioned on observations of Y is denoted as the conditional entropy $\mathrm{H}(F|Y)$, computed as:

$$\mathrm{H}(F|Y) = -\mathbb{E}_Y \mathbb{E}_{F|\mathbf{y}} \log\left(\mathrm{p}_{F|\mathbf{y}}(\mathbf{f}|\mathbf{y})\right) \tag{4-40}$$

Typically, for problems in which variables take on discrete states, a base of 2 is used in the logarithm for defining entropy, whereas in problems with continuous variables, the natural logarithm is used. The marginal and conditional entropy of G, S, or $\mathbb{s}$ can also be defined in a similar way, and used in cases where reduction in uncertainty for these variables directly is of more interest than uncertainty reduction in the underlying random variables F.

The conditional entropy can be used as an objective for sensor placement as follows:

$$\mathrm{M}_{\mathrm{Ent}}(Y) = \frac{\mathrm{H}(F) - \mathrm{H}(F|Y)}{\mathrm{C}(Y)} \tag{4-41}$$

where the same benefit-cost formulation as Equation 4-34 is used here for the entropy reduction (the difference between the marginal and conditional entropies), since entropy is not measured directly in monetary terms. Again, this ratio represents an algorithmic approach to avoid greedy selection of moderately informative but high-cost sensors instead of cheaper sensors which together can provide a greater overall reduction in uncertainty for a lower total cost. However, it can also be used as an overall objective, i.e., for selecting measurement sets which are most informative relative to their cost.

Note the similarity between the numerator of this objective and the VoI as in Equation 4-6; both terms are expressed as the difference between a fixed prior quantity ($\mathbb{E}\mathrm{L}(\emptyset)$ or $\mathrm{H}(F)$) and a posterior quantity that is a set function of Y ($\mathbb{E}\mathrm{L}(Y)$ or $\mathrm{H}(F|Y)$). The former quantities can both be ignored for purposes of minimization, as they are invariant with respect to the chosen measurement set. Furthermore, the differences are both non-negative, due to the non-negativity of the VoI and the fact that $\mathrm{H}(F) \geq \mathrm{H}(F|Y)$ for any Y. It is also interesting to note the similarities in the functional forms of the entropy and the expected loss. For example, comparing Equations 4-1 and 4-39, the latter can be seen as a special case of the former with $\mathrm{L}(\mathbf{f}, \mathbf{a}) = -\log\left(\mathrm{p}_F(\mathbf{f})\right)$. Comparing Equations 4-5 and 4-40, the same relationship can be seen with the substitution $\mathrm{L}(\mathbf{f}, \mathbf{a}) = -\log\left(\mathrm{p}_{F|\mathbf{y}}(\mathbf{f}|\mathbf{y})\right)$; however, in this latter case, the dependence on $\mathbf{y}$ is not captured by the loss function, and so the conditional entropy cannot be completely expressed as a special case of the VoI.

4.4.3.1. Mutual information

The conditional entropy is used indirectly in the work of Krause (2008) to define the mutual information sensing metric that is used for submodular near-optimal sensor placement. Mutual information measures the amount of shared uncertainty between two sets of random variables. In the context of sensor placement, a useful metric to define is the mutual information between a set of selected measurements Y and the remaining set of unselected measurements $\mathcal{Y}\backslash Y$. This mutual information is denoted as:

$$\mathrm{MI}(Y; \mathcal{Y}\backslash Y) = \mathrm{H}(\mathcal{Y}\backslash Y) - \mathrm{H}(\mathcal{Y}\backslash Y|Y) = \mathrm{H}(Y) - \mathrm{H}(Y|\mathcal{Y}\backslash Y) \tag{4-42}$$

The key idea behind the use of mutual information as a sensing metric in this way is the fact that the selected measurements should “cover” the information provided by the set of all possible measurements as much as possible, i.e., the amount of uncertainty in $\mathcal{Y}\backslash Y$ that is not explained by Y should be minimized. An objective for optimal sensing based on this idea is:

$$\mathrm{M}_{\mathrm{MI}}(Y) = \frac{\mathrm{MI}(Y;\mathcal{Y}\backslash Y)}{\mathrm{C}(Y)} \tag{4-43}$$

Again, a benefit-cost formulation is used here to trade off the non-monetized mutual information with the sensing cost. Note that this represents an algorithmic approach to avoid suboptimal greedy algorithm performance using this metric; both the benefit-cost ratio, depicted above, as well as the mutual information directly are used to guide the greedy selection of measurements. The final selected measurement sets via both approaches are compared to determine which of the two placements is to be selected. The reader is referred to the work of Krause (2008) for further details, investigation of the performance of this metric, guarantees on near-optimality of greedily obtained solutions in cases where sensing costs are equal or unequal, and extensions of this basic concept to many different sensor placement applications.

4.4.3.2. Gaussian entropy

The marginal or conditional entropy for a multivariate Gaussian random variable F, as in Equation 3-4, is computable in closed-form as (Cover and Thomas, 2006):

$$\mathrm{H}(F) = \frac{1}{2}\ln(\det(2\pi e\boldsymbol{\Sigma}_F)) \quad \mathrm{H}(F|Y) = \frac{1}{2}\ln\left(\det\left(2\pi e\boldsymbol{\Sigma}_{F|Y}\right)\right) \tag{4-44}$$

where $\ln(\cdot)$ denotes the natural logarithm and $\det(\cdot)$ the matrix determinant, and where e is the natural number. Note that the conditional entropy is not a function of the observation value $\mathbf{y}$, and therefore taking the expectation over possible outcomes of this observation is not necessary for computing the conditional entropy for Gaussian random variables.

4.4.3.3. Binary state entropy

Referring back to the binary component of Section 4.3, the entropy of the component state can also be expressed as a function of the failure probability. This entropy is shown in Figure 4-4. Note that the entropy, as with the misclassification rate in Figure 4-3, takes on its maximum value at $P_f = 50\%$, where the binary state of the component is the most uncertain.

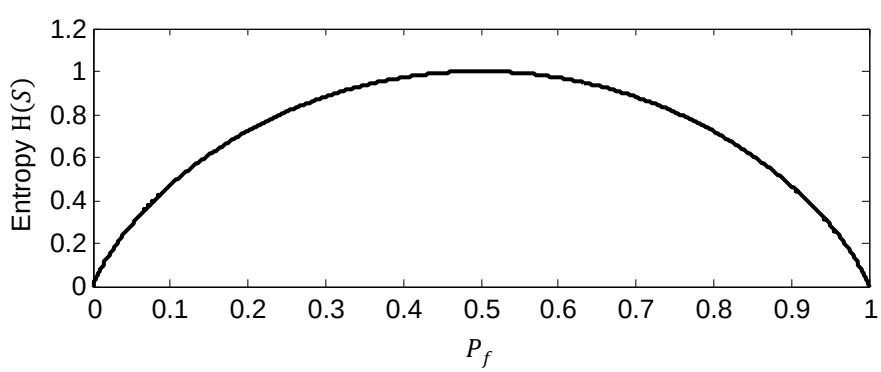

Figure 4-4: Entropy as a function of failure probability for a binary component.

In this case, the entropy of the component can be expressed as the minimum expected loss using a specific definition for the loss function:

$$\mathrm{L}(s,a) = s(-\log(1-a)) + (1-s)(-\log(a)) \tag{4-45}$$

where the action is considered to be a continuous quantity $a \in [0,1]$. Each choice from this infinite set of possible actions defines an expected loss line (as in Figure 4-1a for the two possible actions considered there) which is tangent to the entropy curve of Figure 4-4 at one point. The set of all actions therefore defines an upper-bound to the entropy as a function of failure probability in a binary component. Note that this same formulation applies both to the marginal and conditional entropy. Further note the fact that $\mathrm{H}(S) \geq \mathrm{H}(S|Y)$ is again a consequence of Jensen's Inequality and the concavity of the entropy as a function of failure probability, as was discussed for the expected loss in Section 4.3.

4.4.3.4. Sum of marginal versus joint entropy

Building on the idea of the entropy of binary component state variables discussed above, the idea of entropy evaluation for multiple components is examined here in more detail. When quantifying the uncertainty of a set of variables, there are two potential approaches. First, the joint entropy of the set can be quantified, as in Equations 4-39 and 4-40. However, for multivariate variables, this becomes more difficult, as the expectations must be evaluated over higher-dimensional probability spaces. Second, the entropy of each variable can be computed, and summed. The result is the sum of marginal entropies. In general, this is computationally simpler, since the expectation need only be evaluated with respect to a single variable for each marginal entropy computation. These two approaches are not equivalent; in fact, the sum of marginal entropies for a set of variables is greater than or equal to the joint entropy, where this equality holds if and only if the variables are mutually independent.

The sum of the marginal entropies for component state variables S is expressed as:

$$\mathrm{H}_{\mathrm{marginal}}(S) = \sum_{i=1}^{n} \mathrm{H}(S_i) \geq \mathrm{H}(S) \tag{4-46}$$

An objective for sensor placement based on this sum of marginal entropies, following the formulation of Equation 4-41, is:

$$\mathrm{M}_{\mathrm{Ent(marginal)}}(Y) = \frac{\mathrm{H}_{\mathrm{marginal}}(S) - \mathrm{H}_{\mathrm{marginal}}(S|Y)}{\mathrm{C}(Y)} \tag{4-47}$$

where, as before, the prior entropy $\mathrm{H}_{\mathrm{marginal}}(S)$ is constant with respect to Y, and the conditional entropy $\mathrm{H}_{\mathrm{marginal}}(S|Y)$ is computed as in Equation 4-46, substituting $\mathrm{H}(S_i|Y)$ for $\mathrm{H}(S_i)$. Note also that, although being defined here with respect to the component states S for consistency with the application discussed in Chapter 9, this objective may also be defined using the entropy of F, G, or $\mathbb{S}$ as appropriate.

In the context of sensor placement in large systems, this alternative to the joint entropy has several advantages. As mentioned, it is computationally more tractable than the joint entropy metric; its decomposability across components as shown in Equation 4-46 is analogous to the decomposition of loss in Equation 3-1, and has corresponding implications for the metric's computation, as discussed in Section 5.2. Also, the sum of marginal entropies upper-bounds the joint entropy, and so a reduction in this bound will tend to reduce the joint entropy (although this may not hold true in all cases). Finally, this formulation captures a relevant feature of uncertainty in many problems of interest. For example, consider a system in which an additional component is added, whose performance is identical to that of an existing component. This addition does not increase the joint entropy of S, but does increase the sum of marginal entropies. In terms of management of the system, the addition of a component tends to be considered as a source of uncertainty, even if its behavior is tied to that of another component. Therefore, the sum of marginal entropies may better capture the perceptions and concerns of infrastructure managers.

4.4.3.5. Qualitative comparison

The conditional entropy and VoI metrics (including the prediction error and misclassification rate as special cases) represent alternative approaches to addressing the problem of optimal sensor placement. Conditional entropy is defined directly from the probabilistic model of the system to measure uncertainty in its variables, while VoI uses additional input about available actions and costs for system management. Thus, when supporting decision-making is the ultimate goal of sensing, the VoI metric should be applied, as it will take relevant costs and actions for system management into account when prioritizing which uncertainties to reduce through sensing, while the conditional entropy will not. Furthermore, since VoI and sensing cost are expressed in the same units, they can be compared directly as in Equation 4-12 without the need to resort to a benefit-cost formulation, as in Equation 4-41. This allows the objective to automatically determine the optimal number of sensors to be used, i.e., sensors should be included only when the marginal benefit of adding a sensor to the network is greater than its marginal cost. Additionally, the VoI only assigns value to information that can influence the actions undertaken by a managing agent. That is, if for every observation $\mathbf{y}$ of Y, the optimal posterior action of Equation 4-4 matches the prior action of Equation 4-2, then $\mathrm{VoI}(Y)$ will be zero (Koller and Friedman, 2009). Thus, only information that can potentially be used to improve decision-making is valued.

Application of the VoI metric to sensor placement depends on an agent defining his or her loss (or utility) function for system management. However, some agents may prefer to define their loss in terms of uncertainty; in that case, conditional entropy would best capture their true objective. There are several practical reasons why an agent might choose to define their objective this way. Decision-makers sometimes weigh difficult-to-quantify consequences, where attempts to assign monetary values to these may meet with significant social or political disagreements. Information from the proposed sensing network may also be needed to support many different problems of decision-making for the system, and so selecting a single context in which to define

the VoI is not appropriate. Alternatively, the sensing network might be involved in a study of the system to accurately characterize its behavior, e.g., to determine appropriate modeling parameters W to describe it. For these reasons, infrastructure managers may prefer to apply the conditional entropy metric to guide optimal sensing, as it is insensitive to cost considerations and has the reduction of posterior uncertainty as its aim.

Finally, note that these metrics can be viewed as alternative mappings from a conditional probability distribution over random field variables F given information Y to a quantitative measurement of sensor benefit $\mathrm{M}(Y)$. For the VoI metric, this mapping is informed by the loss function specified by the infrastructure manager, and thus allows uncertainties in different variables to be weighted differently based on the different costs and available management actions for the system. For the conditional entropy metric, uncertainty is penalized equally across all variables, resulting in a measure of system uncertainty that is insensitive to economic or management distinctions between components. The specific goals and concerns of the manager will determine which of these mappings should be applied in any particular problem. This view of the metrics as alternative mappings is extended in Section 5.2.2 to develop a common framework for their evaluation under specific assumptions.

4.5. Non-Submodularity of the Value of Information Metric

As introduced in Section 1.2.4, the set-functional property of submodularity, which is intuitively understood as a diminishing returns property, allows a greedy optimization approach such as that presented in Algorithm 1 to be applied to maximize the set function's value with certain statistical guarantees on the near-optimality of the solution. Specifically, applying a greedy optimization approach with a submodular objective function guarantees that the value of the obtained solution will be within $1 - 1/e$ (about 63%) of the true optimal value that would be attained through an exhaustive search of the solution space (Nemhauser et al., 1978). Thus, submodularity is an important property for providing guarantees about the obtained solution in large combinatorial optimization problems where an exhaustive search is infeasible. This section will discuss the submodular characteristics (or lack thereof) of the VoI metric, as well as some of the alternative metrics considered in Section 4.4. The basic modeling assumptions of Section 3.3 are relaxed in this section.

One of several equivalent definitions of the submodular property is that set function $\mathrm{M}(\cdot)$ is submodular if, for any subsets Y_1 and Y_2 of $\mathcal{Y}$:

$$\mathrm{M}(Y_1 \cup Y_2) - \mathrm{M}(Y_1) \leq \mathrm{M}(Y_2) - \mathrm{M}(Y_1 \cap Y_2) \tag{4-48}$$

This motivates the understanding of submodularity as a diminishing returns property, since the left-hand side, which indicates the additional benefit as measured by the function of adding set $Y_2 \backslash \{Y_1 \cap Y_2\}$ to Y_1, is less than or equal to the right-hand side, which indicates the benefit of

adding the set to $Y_1 \cap Y_2$ only, with $Y_1 \cap Y_2 \subseteq Y_1$. In other words, the benefit of adding to a smaller set is greater than the benefit of adding to a larger one.

It is well-known that VoI is not in general a submodular metric. This non-submodularity can be demonstrated in an intuitive example, motivated by a variant on a game of Blackjack. Two random variables f_1 and f_2 can each take on integer values from 1 to 10 inclusive with equal probability. The player choses an action a, and receives a reward value of 1 (i.e., a loss of -1) if the total of $f_1 + f_2 + a$ equals 21. Without any additional information, the player's prior optimal action is to select $a = 10$, yielding an expected reward of 0.1.

Consider that prior to choosing his or her action, the player has a choice to observe one or both of the random variables. Let $y_1 = f_1$ and $y_2 = f_2$. Since the player will always be able to choose an action that guarantees them the reward after observing both variables, the expected reward is 1, and the VoI for observing y_1 and y_2 is 0.9. If either y_1 or y_2 is observed, the player knows one of the variables, and thus has a one-in-ten chance of successfully guessing the other variable (i.e., of taking an action that will sum to 21 with the other variables) and of receiving the reward. Therefore, the posterior expected loss after observing either y_1 or y_2 is 0.1, and so the VoI of either observation on its own is 0. Note that the posterior optimal action in this case is dependent on the value of the observation, but the posterior expected reward is still the same as in the prior case, taking the prior optimal action.

In this example, the VoI is clearly not submodular, since $\mathrm{VoI}(\{y_1\}) = \mathrm{VoI}(\{y_2\}) = 0$ while $\mathrm{VoI}(\{y_1, y_2\}) = 0.9$, and (by definition) $\mathrm{VoI}(\emptyset) = 0$. Thus, the marginal benefit of any one measurement is 0, but their joint benefit is 0.9; clearly there are increasing returns for the additional measure, rather than decreasing ones.

Nevertheless, there are specific cases where VoI does satisfy the definition for submodularity. For the prediction error loss function (as discussed in Section 4.4.1) in Gaussian process models, under certain assumptions on the model structure, the VoI is submodular (Das and Kempe, 2008; Krause et al., 2008b). Furthermore, in the case of an independent system as in Section 3.2.1 with a decomposable loss function as in Equation 3-1, the VoI for these independent measurements is submodular (a proof is provided in Appendix B).

In terms of other sensor placement objectives, the conditional entropy metric is submodular in several cases of interest, such as when measurements Y are conditionally independent given random field variables F, or when Y describes direct observations of variables in F (see the proof in Appendix B). Furthermore, the mutual information between selected and unselected measurements, as discussed in Section 4.4.3.1, is always submodular (Krause, 2008). These properties can make the conditional entropy and mutual information metrics attractive in applications where robust bounds on the quality of the greedily obtained solution are necessary.

Despite the fact that the VoI metric is not generally submodular, the greedy optimization algorithm remains a valid technique for sensor placement using this objective. In fact, even in problems with submodular objectives, the guarantees given by this property are often in practice

quite loose, with the greedy optimization algorithm providing solutions that are optimal or near-optimal in many practical cases (Krause, 2008). In other words, even though the greedy algorithm may perform arbitrarily poorly when optimizing non-submodular objectives, there are many cases in which this will not occur.

What is indicated for the use of the VoI metric in sensor placement optimization is therefore a consciousness of the limitations of greedy optimization. Practically, this can be maintained through intuitive "sanity checks" on the performance of the algorithm. For example, using the bounds of Equation 4-7, an upper-limit can be established for the VoI of any sensor placement scheme. Comparing the value of the greedily optimized scheme to this upper bound can provide a heuristic for determining if the resulting placement is "good enough", by comparing to what can be achieved in the limit of complete information, or identifying when the algorithm is performing poorly, if the value is too far from this bound. Furthermore, testing the performance of the greedy approximation algorithm against that of an exact solution approach in a tractably small system can identify any serious problems that might be faced when the approach is scaled up. Finally, a basic understanding of the problem can be used to recognize situations that might lead the greedy approach to perform poorly, such as when the benefits of additional information are clearly increasing rather than diminishing.

Throughout the remainder of this work, and especially in the case studies presented in Chapter 8, Chapter 9, and Chapter 10, applications of problem-specific intuition as well as empirical evidence will be used to judge the performance of the greedy optimization approach to VoI maximization. Section 7.2.3 also presents a specific example of obviously suboptimal performance by the algorithm and an intuitive countermeasure. Furthermore, Section 11.1 presents further empirical evidence to support the use of the greedy optimization approach for maximizing VoI.

Chapter 5

Metric Evaluation and System Topology

The use of the VoI metric for sensor placement can be computationally demanding in general. The source and degree of this complexity is discussed in Section 5.1. This chapter focuses on presenting computational approaches for the efficient evaluation of VoI in different cases. In particular, when the loss function is decomposable, as in Equation 3-1, the VoI metric can be decomposed in an analogous manner, leading to increased computational efficiency. Section 5.2 discusses this general decomposition, as well as further computational efficiency gains available when Gaussian process models for the random fields and binary component state models are used, following the assumptions of Section 3.3. Section 5.3 presents a generally applicable sampling-based method for VoI evaluation for cases in which the VoI metric is not decomposable, and discusses potential issues with this approach, including approximate methods for use in series and parallel systems. Finally, Section 5.4 presents a parametric analysis comparing the behavior of the VoI in a five-component system with a decomposable loss function with its performance in two systems with non-decomposable loss functions, namely a parallel and series system with the same number of components. This analysis illustrates some interesting characteristics of the VoI metric in these non-decomposable cases, which might be exploited to allow for efficient approximate sensor placement heuristics to be developed for series systems.

5.1. Computational Difficulties

The evaluation of the VoI, as illustrated in Equation 4-5 for the posterior expected loss, requires several steps. First, the inner expectation over posterior distribution $\mathrm{p}_{F|\mathbf{y}}$ of the loss function must be computed. This requires that the probability $\mathrm{p}_{F|\mathbf{y}}$ be evaluated for each combination of $\mathbf{f}$ and $\mathbf{y}$, which can be difficult if p_F is not a conjugate prior of $\mathrm{p}_{Y|\mathbf{f}}$, as discussed in Section 2.1.2.1. The expectation must then be evaluated over this posterior. In general, the expectation will not have a closed-form expression, so approximate techniques such as sampling-based methods must be applied. MCMC methods represent a well-developed class of general techniques for sampling from the posterior distribution for the purpose of approximating posterior expectations, and can be applied to evaluate these expectations (e.g., Au and Beck, 2001; Papaioannou et al., 2015).

Next, the minimum conditional expected loss over all possible actions must be evaluated for each potential observation. For actions on a continuous space, this optimization problem may be convex (see e.g., Boyd and Vandenberghe, 2004) or non-convex (see e.g., Brandimarte, 2006). For discrete actions, this problem is again one of combinatorial optimization, and the only guarantee of an optimal solution in all cases is an exhaustive search through all possible actions, which is computationally expensive. A greedy search algorithm can again be applied to this problem, similar to the approach of Algorithm 1 (except for minimizing over actions, rather than maximizing over observation sets). However, as with the sensor placement optimization problem, there are in general no guarantees on the optimality of the solution, or even its near-optimality if the submodularity of the conditional expected loss cannot be established.

Finally, an outer expectation over potential observations of Y is needed. In general, approximate methods, e.g., sampling-based methods such as Monte Carlo sampling, can also be applied to evaluate this outer expectation, as above for the inner expectation. As all the individual steps of posterior expected loss evaluation can be computationally intensive, and several of these are repeated in the prior expected loss computation, the overall evaluation of VoI is a challenging problem in general. Formally, in terms of computational complexity classes, the problem of evaluating VoI in Naïve Bayes models (a relatively simple class of Bayesian Network PGMs with a tree structure, known as a polytree) is #**P**-complete, and the problem of optimizing measurements based on VoI is $\mathbf{NP}^{\mathbf{PP}}$-hard for polytrees (Krause and Guestrin, 2009).

To quantify this challenge in a specific case, consider that a system with n binary components is to be monitored via a set Y of n_Y binary observations, and managed with binary management action choices, with one choice associated with each component. Therefore, there are 2^{n_Y} possible outcomes of the observation set, 2^n possible joint states for the components, and 2^n possible management actions that can be taken across the system. The loss function in this case is expressed in the form of $\mathrm{L}(\mathbf{s}, \mathbf{a})$, and is not decomposable in general. Complexity of the VoI evaluation can be quantified in terms of the complexity for evaluating the posterior expected loss, since this is a necessary step in evaluating the VoI and one that must be repeated for every potential measurement set during sensor placement optimization, whereas the prior expected loss need only be computed once. Computation of the posterior expected loss requires that, for each potential measurement outcome, the posterior probabilities of each possible joint state be evaluated, and the expected loss taken across these probabilities and minimized across all potential actions.

The complexity of this process can be quantified using order notation, where for example $\mathrm{O}(n^2)$ indicates a process whose complexity increases with the square of n. The complexity of each step of the evaluation of the posterior expected loss can thus be quantified, and the upper limiting order of this complexity identified. Note that here, the size of the search space (i.e., the number of joint variable states and actions which must be checked) is being quantified. This may not correspond exactly to the algorithmic complexity, as implemented algorithms for solving particular problems can have lower complexity than the search space size. In particular, this

chapter will examine how complexity can be reduced by exploiting the problem structure in some specific cases. What is presented now is therefore the worst-case complexity in the most general systems, requiring exhaustive search through all possibilities.

First, the posterior probability distribution $\mathrm{p}_{S|\mathbf{y}}$ must be evaluated for each combination of $\mathbf{s}$ and $\mathbf{y}$, requiring $\mathrm{O}(2^n 2^{n_Y})$ evaluations. The loss function must be computed for each combination of $\mathbf{s}$ and $\mathbf{a}$, requiring $\mathrm{O}(2^n 2^n)$ evaluations. Note that these loss values will not change for different observations $\mathbf{y}$ of Y, and so may be precomputed and stored, although the storage requirements may be prohibitive in larger problems. Evaluation of the conditional expected loss is needed for each of 2^{n_Y} observations and 2^n actions, each requiring a sum over a product of 2^n losses and state probabilities. Overall, this represents $\mathrm{O}(2^{n_Y} 2^n 2^n)$ operations. Minimization over actions must be performed for each of 2^{n_Y} observations, and requires a minimization over 2^n action choices. Complexity of finding the minimum of an unsorted set is linear (Newman, 2010), and so the minimization steps require a total of $\mathrm{O}(2^{n_Y} 2^n)$ operations. Finally, evaluating the outer expectation requires evaluating the observation probability distribution p_Y 2^{n_Y} times and the summing over 2^{n_Y} products, for an overall complexity of $\mathrm{O}(2^{n_Y})$. The bottleneck of this process is the evaluation of the inner expectation, with a computational complexity of $\mathrm{O}(2^{n_Y} 4^n)$.

For reference, the complexity for computing the posterior conditional entropy $\mathrm{H}(S|Y)$ in this problem would be $\mathrm{O}(2^{n_Y} 2^n)$, i.e., the posterior probabilities of all states and observations would need to be evaluated, and the complexity for evaluating the sum of marginal entropies $\mathrm{H}_{\mathrm{marginal}}(S|Y)$ as in Equation 4-46 would be $\mathrm{O}(n 2^{n_Y})$, i.e., the marginal state probability for each component would need to be evaluated for every potential observation.

An additional level of complexity is added by the optimization problem. For exact optimization of Y from a set $\mathcal{Y}$, there are $2^{n_\mathcal{Y}}$ possible sets, as discussed in Section 2.2.1. Thus, exact sensor placement optimization using the VoI metric in this problem has complexity $\mathrm{O}(2^{n_\mathcal{Y}} 2^{n_Y} 4^n)$, with $n_Y \leq n_\mathcal{Y}$. Adopting the greedy heuristic reduces this to $\mathrm{O}\left(n_\mathcal{Y}^2 2^{n_Y} 4^n\right)$. As can be seen, even when adopting a greedy optimization approach, the complexity of sensor placement using the VoI metric grows exponentially with the number of system components and the number of measurements taken. When dealing with continuous rather than discrete variables, the complexity of the problem cannot be so easily quantified; however, analogous problems with the exponential growth of the solution space dimensionality occur.

Over the following sections, it is shown that in certain special cases of the system topology and choice of the loss function, in particular when the loss function is decomposable across system components, this computational complexity can be reduced to be linear in the number of components. Additionally, in some cases where the loss function is not decomposable, approximate techniques can be used to reduce the computational complexity. Since the choice of loss function impacts the computational complexity of VoI evaluation, the selection of an appropriate loss function can trade off the need to accurately capture the economic cost of a given problem to an infrastructure manager against the computational cost of VoI evaluation.

Significant gains in computational efficiency might then motivate a managing agent to use an approximate loss function of a simpler problem class (e.g., a decomposable function) to substitute for a more accurate but more computationally demanding loss function. Chapter 6 presents several examples where a computationally efficient loss function serves as a reasonable proxy to another, more intensive loss function.

5.2. Decomposable Loss Functions

The assumption of a decomposable loss function, or alternatively a cumulative system topology as discussed in Section 3.2.2, allows the VoI metric to be decomposed across the components of the system, reducing the computational effort associated with its evaluation. This decomposition is motivated by concepts in dynamic programming (e.g., Bertelè and Brioschi, 1972).

To derive this decomposition of the VoI metric, begin by substituting Equation 3-1 into Equation 4-5:

$$\mathbb{E}\mathrm{L}(Y) = \mathbb{E}_Y \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}} \sum_{i=1}^{n} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) \tag{5-1}$$

First, making use of the linearity of the expectation, the expectation of the sum can be written as the sum of expectations as follows:

$$\mathbb{E}\mathrm{L}(Y) = \mathbb{E}_Y \min_{\mathbf{a}\in\mathcal{A}} \sum_{i=1}^{n} \mathbb{E}_{F_i|\mathbf{y}} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) \tag{5-2}$$

where, since each component-level loss function is only a function of the subset F_i of F, the expected loss over all F is equivalent to an expected loss over each subset.

Next, note that the component-level loss is a function only of actions A_i, rather than of all actions A. Furthermore, the minimum of a summation can be achieved by minimizing each summand when these are functions of disjoint action sets. Therefore, the minimum of the total conditional expected loss across components is equivalent to the sum of the minimum of the conditional expected losses for each component:

$$\mathbb{E}\mathrm{L}(Y) = \mathbb{E}_Y \sum_{i=1}^{n} \min_{\mathbf{a}_i\in\mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) \tag{5-3}$$

Finally, again using the linearity of the expectation, the outer expectation can be moved inside the summation:

$$\mathbb{E}\mathrm{L}(Y) = \sum_{i=1}^{n} \mathbb{E}_Y \min_{\mathbf{a}_i\in\mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) = \sum_{i=1}^{n} \mathbb{E}\mathrm{L}_i(Y) \tag{5-4}$$

Note that this expectation is still performed over observations of all measures in Y. If this cumulative system were also an independent system, as discussed in Section 3.2.1, this expectation could also be expressed as an expectation over Y_i, since the expected loss for this component would be independent of any measurements associated with other components.

In the case where the system-level loss function is decomposable, the system-level VoI metric can therefore be decomposed across system components as:

$$\mathrm{VoI}(Y) = \sum_{i=1}^{n} \mathrm{VoI}_i(Y) \tag{5-5}$$

where:

$$\mathrm{VoI}_i(Y) = \mathbb{E}\mathrm{L}_i(\emptyset) - \mathbb{E}\mathrm{L}_i(Y) \tag{5-6}$$

Note here that the prior expected loss will decompose in the same way as the posterior, shown in Equation 5-4, only without conditioning on Y.

In terms of computational complexity, the decomposition of the VoI metric reduces the computational cost growth to be linear in the number of components, rather than exponential. Revisiting the problem of a system with binary component states, actions, and observations from Section 5.1, the computational bottleneck is now the evaluation, for each of n components, of the conditional expected loss given each of 2^{n_Y} observations, 2 component actions, and 2 component states, representing $\mathrm{O}(n2^{n_Y})$ operations. Thus, the growth of complexity is linear in the number of components rather than exponential (assuming of course that the number of component states and management actions does not grow with the size of the system).

Other sensor placement objectives benefit from decomposability as well. Recall from Section 4.4.3.4 that the sum of marginal entropies can be used as a decomposable alternative to the joint entropy. Thus, in a problem with binary observations and states, the computational complexity of evaluating the sum of marginal entropies is $\mathrm{O}(n2^{n_Y})$, in contrast to $\mathrm{O}(2^n 2^{n_Y})$ for the joint entropy.

Two specific decomposable loss functions that allow for additional computational savings are examined in further detail in the following subsections.

5.2.1. Efficient Evaluation for Gaussian Prediction Error

In the case of the field prediction error loss function, as discussed in Section 4.4.1, under the assumption of a Gaussian random field and observations, as discussed in Section 3.3, the VoI can be easily assessed in closed-form. Beginning with Equation 4-32, note that the prediction error loss function is decomposable, with $n = n_X$, i.e., the random field variables constitute the "components" of this system. Substituting this loss function into the posterior expected loss of Equation 5-4 gives:

$$\mathbb{E}\mathrm{L}_{\mathrm{pred}}(Y) = \sum_{i=1}^{n_X} \mathbb{E}_Y \min_{a_i} \mathbb{E}_{f_i|\mathbf{y}}(f_i - a_i)^2 \tag{5-7}$$

where a_i is the continuous estimate of f_i. Now note that the estimate that will minimize the expected prediction error of the Gaussian random varaible f_i conditional to measuremnet $\mathbf{y}$ is the posterior mean of f_i conditioned on this measurment:

$$\text{argmin}_{a_i} \mathbb{E}_{f_i|\mathbf{y}}(f_i - a_i)^2 = \mu_{f_i|\mathbf{y}} \tag{5-8}$$

where this posterior mean value can be computed in closed-form as the ith element of posterior random field mean vector $\boldsymbol{\mu}_{F|\mathbf{y}}$, as in Equation 3-10. Substituting this value into the posterior expected loss yields:

$$\mathbb{E}\text{L}_{\text{pred}}(Y) = \sum_{i=1}^{n_X} \mathbb{E}_{f_i,Y}\left(f_i - \mu_{f_i|\mathbf{y}}\right)^2 \tag{5-9}$$

Now, note that the expected squared difference between the random field and its posterior mean is the posterior variance of the variable:

$$\mathbb{E}_{f_i,Y}\left(f_i - \mu_{f_i|\mathbf{y}}\right)^2 = \sigma_{f_i|Y}^2 \tag{5-10}$$

Substituting this relationship into the posterior expected loss expression yields:

$$\mathbb{E}\text{L}_{\text{pred}}(Y) = \sum_{i=1}^{n_X} \sigma_{f_i|Y}^2 = \text{tr}\left(\boldsymbol{\Sigma}_{F|Y}\right) \tag{5-11}$$

which can be computed using the matrix trace of posterior covariance matrix $\boldsymbol{\Sigma}_{F|Y}$, obtained in closed-form from Equation 3-10. Note that the prior expected loss can likewise be obtained by substituting $\boldsymbol{\Sigma}_F$ for $\boldsymbol{\Sigma}_{F|Y}$ in this expression. Thus, the VoI for prediction error is evaluated in closed-form as:

$$\text{VoI} = \text{tr}(\boldsymbol{\Sigma}_F) - \text{tr}\left(\boldsymbol{\Sigma}_{F|Y}\right) \tag{5-12}$$

When the loss is expressed as a function of the limit-state variables G as in Equation 4-33, the VoI becomes, following the relationship of Equation 3-13:

$$\text{VoI} = \text{tr}(\boldsymbol{\Sigma}_G) - \text{tr}\left(\boldsymbol{\Sigma}_{G|Y}\right) = \text{tr}\left(\boldsymbol{\Xi}_G \boldsymbol{\Sigma}_F \boldsymbol{\Xi}_G^{\text{T}}\right) - \text{tr}\left(\boldsymbol{\Xi}_G \boldsymbol{\Sigma}_{F|Y} \boldsymbol{\Xi}_G^{\text{T}}\right) \tag{5-13}$$

Similarly, for the weighted prediction error using the loss function of Equation 4-35, VoI is:

$$\text{VoI} = \text{tr}(\mathbf{Z}\boldsymbol{\Sigma}_F\mathbf{Z}^{\text{T}}) - \text{tr}\left(\mathbf{Z}\boldsymbol{\Sigma}_{F|Y}\mathbf{Z}^{\text{T}}\right) \tag{5-14}$$

Note here that for $\mathbf{Z} = \boldsymbol{\Xi}_G$, the weighted prediction loss function and the loss function based on prediction of the limit-state variables will lead to the same VoI, and thus the same optimal sensor placements under this objective. This supports the equivalence of these two sensor placement approaches, as alluded to in Section 4.4.1.1.

Furthermore, note that the conditional entropy of random field variables F is also expressible as a function of posterior covariance $\boldsymbol{\Sigma}_{F|Y}$, as in Equation 4-44. This means that the VoI metric will have the same computational complexity as the conditional entropy metric, in the case of a Gaussian random field and the prediction error loss function. This complexity is driven by the computation of $\boldsymbol{\Sigma}_{F|Y}$ via Equation 3-10, which is governed by the inversion of the n_Y by n_Y

matrix $\mathbf{\Sigma}_Y$, leading to an overall complexity of $\mathrm{O}(n_Y^3)$ for this step (Rasmussen and Williams, 2006). In general, n_Y is less than n_X, as only a subset of the discrete locations are usually observed. In terms of the relationship of complexity to n_X, this can be expressed in terms of the need for matrix addition to evaluate the posterior covariance in Equation 3-10, which is $\mathrm{O}(n_X^2)$, or the need to evaluate the matrix determinant in Equation 4-44, which is $\mathrm{O}(n_X^3)$. In either case, the rate of increase is less than the exponential growth exhibited by general VoI evaluation.

5.2.2. Efficient Evaluation for Binary Components

In cases where components have binary states, as outlined in Section 3.3.4, and these states are functions of underlying Gaussian random variables, the exponential dependence of the complexity of VoI evaluation on the number of observations can be reduced. First, using the formula for the posterior mean of the random field given a set of observations, as in Equation 3-10, and the prior distribution of these observations of Equation 3-7, note that the posterior mean vector $\boldsymbol{\mu}_{F|\mathbf{y}}$ is an affine function of Gaussian random variable $\mathbf{y}$, and so is itself a Gaussian random variable:

$$\boldsymbol{\mu}_{F|\mathbf{y}} \sim \mathcal{N}\left(\boldsymbol{\mu}_F, \mathbf{\Sigma}_{FY}{\mathbf{\Sigma}_Y}^{-1}\mathbf{\Sigma}_{FY}^{\mathrm{T}}\right) \tag{5-15}$$

where the covaraiance matrix is derived by noting that $\left(\mathbf{\Sigma}_{FY}{\mathbf{\Sigma}_Y}^{-1}\right)\mathbf{\Sigma}_Y\left(\mathbf{\Sigma}_{FY}{\mathbf{\Sigma}_Y}^{-1}\right)^{\mathrm{T}} = \mathbf{\Sigma}_{FY}{\mathbf{\Sigma}_Y}^{-1}\mathbf{\Sigma}_{FY}^{\mathrm{T}}$. Using the relationship between limit-state variables G and random field variables F of Equation 3-11, the posterior distribution parrameters for G are:

$$\boldsymbol{\mu}_{G|\mathbf{y}} = \mathbf{\Xi}_G\boldsymbol{\mu}_{F|\mathbf{y}} + \mathbf{g}_0 \quad \mathbf{\Sigma}_{G|Y} = \mathbf{\Xi}_G\mathbf{\Sigma}_{F|Y}\mathbf{\Xi}_G^{\mathrm{T}} \tag{5-16}$$

Since $\boldsymbol{\mu}_{G|\mathbf{y}}$ is an affine function of Gaussian variable $\boldsymbol{\mu}_{F|\mathbf{y}}$, it is also Gaussian with distribution:

$$\boldsymbol{\mu}_{G|\mathbf{y}} \sim \mathcal{N}\left(\mathbf{\Xi}_G\boldsymbol{\mu}_F + \mathbf{g}_0, \mathbf{\Xi}_G\mathbf{\Sigma}_{FY}{\mathbf{\Sigma}_Y}^{-1}\mathbf{\Sigma}_{FY}^{\mathrm{T}}\mathbf{\Xi}_G^{\mathrm{T}}\right) \tag{5-17}$$

where the mean vector derives from Equation 3-11 and the relationship $\mathbb{E}_Y\boldsymbol{\mu}_{F|\mathbf{y}} = \boldsymbol{\mu}_F$.

Recalling the definition of the reliability index from Equation 3-17, the prior or posterior failure probability for a binary component can be derived from the parameters of the associated limit-state variable as:

$$P_{f,i} = \Phi(-\beta_i) = \Phi\left(-\frac{\mu_{g_i}}{\sigma_{g_i}}\right) \quad P_{f,i|\mathbf{y}} = \Phi\left(-\beta_i(\mathbf{y})\right) = \Phi\left(-\frac{\mu_{g_i|\mathbf{y}}}{\sigma_{g_i|Y}}\right) \tag{5-18}$$

Motivated by these relationships, the reliability vector $\boldsymbol{\beta}$ concatenates the reliability indices of all binary system components. In the prior case or conditional to observation $\mathbf{y}$, this vector is:

$$\boldsymbol{\beta}(\emptyset) = \mathrm{diag}(\mathbf{\Sigma}_G)^{-\frac{1}{2}}\boldsymbol{\mu}_G \quad \boldsymbol{\beta}(\mathbf{y}) = \mathrm{diag}\left(\mathbf{\Sigma}_{G|Y}\right)^{-\frac{1}{2}}\boldsymbol{\mu}_{G|\mathbf{y}} \tag{5-19}$$

where $\mathrm{diag}(\cdot)$ denotes the diagonalization of a matrix, i.e., all off-diagonal elements of the matrix are set to 0. In the prior case, this reliability vector is a constant, and in the posterior case it is a linear function of observation vector $\mathbf{y}$ for a given set of measures Y. Therefore, in the posterior case, the reliability index vector can be treated as a Gaussian random vector:

$$\boldsymbol{\beta} \sim \mathcal{N}\left(\boldsymbol{\mu}_{\beta|Y}, \boldsymbol{\Sigma}_{\beta|Y}\right) \tag{5-20}$$

where:

$$\boldsymbol{\mu}_{\beta|Y} = \mathrm{diag}\left(\boldsymbol{\Sigma}_{G|Y}\right)^{-\frac{1}{2}}\boldsymbol{\mu}_G \quad \boldsymbol{\Sigma}_{\beta|Y} = \mathrm{diag}\left(\boldsymbol{\Sigma}_{G|Y}\right)^{-\frac{1}{2}}\boldsymbol{\Xi}_G\boldsymbol{\Sigma}_{FY}\boldsymbol{\Sigma}_Y{}^{-1}\boldsymbol{\Sigma}_{FY}^{\mathrm{T}}\boldsymbol{\Xi}_G^{\mathrm{T}}\mathrm{diag}\left(\boldsymbol{\Sigma}_{G|Y}\right)^{-\frac{1}{2}} \tag{5-21}$$

Note that the conditioning on observation $\mathbf{y}$ has been dropped for notational convenience, since in the prior case, i.e., if $Y = \emptyset$, then $\boldsymbol{\Sigma}_{\beta|Y} = \mathbf{0}$ and $\boldsymbol{\Sigma}_{G|Y} = \boldsymbol{\Sigma}_G$, and so $\boldsymbol{\mu}_{\beta|\emptyset} = \mathrm{diag}(\boldsymbol{\Sigma}_G)^{-\frac{1}{2}}\boldsymbol{\mu}_G = \boldsymbol{\beta}(\emptyset)$, a constant. Also note that the distribution parameters of $\boldsymbol{\beta}$ are functions of the measurement set Y only, rather than the values $\mathbf{y}$ of these measurements.

Recall from Section 4.3 that, for a single binary component, the minimum expected loss can be expressed as a function of component failure probability, as demonstrated in Figure 4-1a. Component-level mapping function $\mathrm{m}_i(\cdot)$ captures this relationship for the ith component. For instance, in the posterior case, the following relationship holds:

$$\mathbb{E}\mathrm{L}_i(\mathbf{y}) = \min_{a_i \in \mathcal{A}_i} \mathbb{E}_{S_i|\mathbf{y}}\mathrm{L}_i(s_i, a_i) = \mathrm{m}_i\left(P_{f,i|\mathbf{y}}\right) = \mathrm{m}_i\left(\Phi\left(-\beta_i(\mathbf{y})\right)\right) \tag{5-22}$$

where, using the component reliability index, this function maps an observation to an expected loss, representing an alternative means of computing the component-level conditional expected loss of Equation 5-4 in the case of a Gaussian system with binary components.

The last key step is to note that, conditional to a set of observations Y, the reliability index for the ith component is distributed according to a univariate Gaussian distribution:

$$\beta_i \sim \mathcal{N}\left(\mu_{\beta_i|Y}, \sigma^2_{\beta_i|Y}\right) \tag{5-23}$$

where the parameters of this distribution are the corresponding elements of the mean vector and along the diagonal of the covariance matrix given in Equation 5-21.

Finally, recalling the definition for the posterior expected loss under the decomposability assumption of Equation 5-4 and substituting in Equation 5-22 yields:

$$\mathbb{E}\mathrm{L}(Y) = \sum_{i=1}^{n} \mathbb{E}_Y \mathrm{m}_i\left(P_{f,i|\mathbf{y}}\right) = \sum_{i=1}^{n} \mathbb{E}_{\beta_i|Y}\mathrm{m}_i\left(\Phi(-\beta_i)\right) \tag{5-24}$$

Note here that the multivariate expectation over Y has been transformed into a univariate expectation over β_i using the relationships developed in this section. The evaluation of this expectation over a lower-dimensional space is computationally simpler. Although the binary

complexity quantification of Section 5.1 no longer applies to this case, as Gaussian variables have been explicitly assumed, the dependence of complexity is now governed only by the complexity of computing the distribution parameters for β_i given Y, which no longer grows exponentially in n_Y. The univariate expectation over the distribution for β_i in Equation 5-23 can be evaluated by any applicable means, e.g., using sampling or numerical integration techniques.

Other decomposable metrics associated with binary states can be evaluated in an analogous way. For example, in the case of component misclassification rates of Section 4.4.2 or the sum of marginal component state entropies of Sections 4.4.3.4, similar mappings $\mathrm{m}_i(\cdot)$ from the probability of component failure to the objective value can be defined, as shown in Figure 4-3 and Figure 4-4. The methods for evaluating misclassification rate and sum of marginal entropy in these cases are exactly as illustrated in this section for the VoI.

5.3. Non-Decomposable Loss Functions

For arbitrary definitions of the loss function, where the function is not decomposable as in Equation 3-1, or where the definition of the system contradicts the assumptions of Section 3.3, the techniques described in the previous section will not be applicable. For systems governed by discrete random variables, or where continuous random variables can be feasibly replaced with discrete approximations, exact evaluation of VoI is possible by enumerating all combinations of observation and state variables and actions, and computing the expected losses numerically. However, as discussed in Section 5.1 (using the example of binary variables and actions), the computational requirements of this approach grow exponentially with the number of variables and observations, and so this approach is infeasible for larger systems.

As a more efficient alternative, sampling-based techniques can be used to approximate the VoI. One such technique, based on crude Monte Carlo sampling, is presented in Section 5.3.1. This technique is broadly applicable to all cases of VoI computation, including when the basic modeling assumptions of Section 3.3 are relaxed. It is employed in this work for evaluating VoI in systems with non-decomposable loss functions. For two specific cases, i.e., for systems with series or parallel network topologies, approximate techniques for efficiently optimizing the choice of posterior management actions are discussed in Section 5.3.2.

5.3.1. General Metric Evaluation via Sampling

This section presents a Monte Carlo sampling approach to evaluating the posterior expected loss of Equation 4-5. An equivalent method for evaluating the prior expected loss is possible by removing any conditioning on observations in the steps outlined in this section. First, to approximate the outer expectation of Equation 4-5, one can draw n_{sim} simulated observations from the distribution over observations p_Y to approximate the expectation as:

$$\mathbb{E}_Y \mathbb{E}\mathrm{L}(\mathbf{y}) \approx \frac{1}{n_{\mathrm{sim}}} \sum_{h=1}^{n_{\mathrm{sim}}} \mathbb{E}\mathrm{L}\left(\mathbf{y}^{(h)}\right) \quad \left\{\mathbf{y}^{(h)}\right\}_{h=1}^{n_{sim}} \sim \mathrm{p}_Y \tag{5-25}$$

This approximate approach reduces the complexity of evaluating the expectation, e.g., in the case of discrete variables, from an exponential growth in n_Y to a linear growth in n_{sim} (ignoring potential issues related to drawing multivariate samples of $\mathbf{y}$).

To approximately evaluate the inner expectation of Equation 4-3 given simulated observation $\mathbf{y}^{(h)}$, one can draw a number n'_{sim} of joint random field variable values from the posterior distribution $\mathrm{p}_{F|\mathbf{y}^{(h)}}$, obtaining:

$$\mathbb{E}_{F|\mathbf{y}^{(h)}}\mathrm{L}(\mathbf{f},\mathbf{a}) \approx \frac{1}{n'_{\text{sim}}}\sum_{j=1}^{n'_{\text{sim}}} \mathrm{L}\left(\mathbf{f}^{(j)},\mathbf{a}\right) \quad \left\{\mathbf{f}^{(j)}\right\}_{j=1}^{n'_{\text{sim}}} \sim \mathrm{p}_{F|\mathbf{y}^{(h)}} \tag{5-26}$$

This approach reduces the evaluation complexity from exponential growth in n_X to linear growth in n'_{sim} (again ignoring potential issues related to drawing samples from the conditional distribution).

Methods for determining appropriate values of n_{sim} and n'_{sim} and for quantifying variance, bias, and confidence intervals in the VoI approximations resulting from this approach are developed by Oakley et al. (2010). Note that those methods have been developed in the context of perfect observations, but apply to VoI evaluation generally. Bias in the estimation of the VoI derives from the minimization over actions for a finite number of sampled states; although the Monte Carlo approximation technique in Equation 5-26 provides an unbiased estimate for $\mathbb{E}_{F|\mathbf{y}^{(h)}}\mathrm{L}(\mathbf{f},\mathbf{a})$, minimization over these estimates introduces a bias into the approximation of Equation 5-25, potentially resulting in underestimation of the loss, and therefore in overestimation of the VoI (Oakley et al., 2010).

To perform minimization over actions in evaluating Equation 4-3 in discrete problems, an exhaustive search through possible actions is generally necessary. As discussed in Section 5.1, this can be computationally prohibitive, and so approximate approaches to minimization may be applied. One such approach is to perform an exhaustive search over a restricted subset $\tilde{\mathcal{A}} \subset \mathcal{A}$ of possible management actions. Methods for generating restricted subsets are presented in the following subsection for series and parallel systems with binary component states and management actions.

5.3.2. Approximate Techniques for Series and Parallel Systems

This subsection develops efficient approximate techniques for minimizing over the choice of management actions in series and parallel systems. In this subsection, the modeling assumptions of Section 3.3 are applied. Furthermore, the prototypical loss function for a non-cumulative system described in Section 4.2.2 and formulated in Equation 4-17 is applied.

In the case of a parallel system, one need only consider two potential actions for managing the system: do nothing for every component, or repair the component with the lowest repair cost. Since only one component needs to be operational or repaired to avoid the failure cost, there is

no need to consider additional actions. Therefore, the number of potential actions that need be considered is reduced from 2^n to 2. Note that this reduced action set must contain the optimal action, and so this technique is exact for parallel systems.

In the case of a series system, under the assumption that the repair costs are equal for all components, a reduced set of management actions of size $n+1$ can be considered. This is motivated by the fact that, in a series system where the binary component states are independent, the prior probability of system failure is:

$$\mathbb{P}(\mathbb{S}'_{\text{series}}(\mathbf{s},\mathbf{a})=0)=1-\prod_{i:a_i=0}\left(1-P_{f,i}\right) \tag{5-27}$$

where the notation on the product indicates that only components that are not repaired are included. Based on this relationship, note that for any set of n_r components to be repaired, the prior probability of system failure will be smallest if the n_r components with the highest marginal probabilities of failure are selected for repair.

While this management strategy is optimal when component failures are independent, it can be assumed that the same strategy will provide a reasonable approximation to the optimal strategy even where failures are non-independent. This assumption thus motivates a heuristic for the optimization of management actions for series systems as follows. First, the probabilities of the component failures are computed as $P_{f,i}$ in the prior case and $P_{f,i|\mathbf{y}}$ in the posterior case for the ith component. Components are then ranked according to this failure probability from highest to lowest, with $i(n_r)$ denoting the index value of the component with the n_r^{th} highest failure probability. Potential management actions are then defined, where $\mathrm{a}^{(n_r)}$ indicates the vector of management actions such that the n_r components with the highest failure probabilities are selected for repair, or:

$$\mathbf{a}^{(n_r)}=\left\{a_{i(1)}=1,\dots,a_{i(n_r)}=1,a_{i(n_r+1)}=0,\dots,a_{i(n)}=0\right\} \tag{5-28}$$

The reduced subset $\tilde{\mathcal{A}}$ of management actions to be considered is then $\left\{\mathbf{a}^{(0)},\dots\mathbf{a}^{(n)}\right\}$, and thus the number of potential joint management actions that must be evaluated when computing VoI is reduced from 2^n to $n+1$. Note that the resulting management action set chosen will be only an approximation to the true optimal management action set in general.

The above approach represents a greedy selection process where components are selected for repair in a specific sequence (although all repairs are to be implemented simultaneously, as has been assumed throughout). It should be noted that in the case of independent components with equal repair costs, as in the example which motivated this approach, the problem of identifying the components to repair can be characterized as a submodular optimization problem. In this case, the log-reliability of the system (where the reliability is one minus the failure probability) is the objective to be maximized through repair actions. This objective is non-decreasing, as performing repairs does not decrease system reliability. Furthermore, this objective is submodular, and in particular it is modular, i.e., the inequality of Equation 4-48 holds in the

equality. This can be seen by noting that for a series system with independent components the log-reliability is the sum of component log-reliabilities for non-repaired components. Repairing a component increases the log-reliability by the same amount regardless of the number of other components which have been repaired. For this reason, the greedy selection process and the resulting approach described above provides guarantees on near-optimality of the solution when component behaviors are independent and when component repairs have the same cost. This observation is motivated by the work of Heckerman et al. (1993).

In the case that component repair costs are unequal, this approximation may perform badly. For example, a large number of components might be repaired at a small cost, significantly reducing the failure probability for the system, while a single repair performed at great cost for the least reliable component may not be able to greatly reduce this failure probability. In such a case, a modified greedy approach to choosing management activities can be adopted, e.g., starting with the do-nothing alternative for all components and choosing one component at a time to repair based on minimizing the resulting expected loss. Note that this greedy selection should be continued until all components have been repaired, as it may be the case that this is the optimal management action, i.e., after all $n+1$ optimal action sets have been chosen (each with a different numbers of components repaired) the final expected losses for all action sets should be compared to determine the optimal one. Finally, note that the conditional expected cost is unlikely to be submodular (intuitively, in a series system, many components may need to be repaired before the system begins to operates, and so the effects of each individual repair may be negligible) and so a greedy approach may perform arbitrarily badly.

5.4. Investigations in Five Component Systems

This section presents a comparative analysis of the VoI metric in multicomponent systems, under different topology assumptions. The purpose of this analysis is to identify similarities in VoI metric values for systems with cumulative topologies, where the metric is relatively efficient to compute, and systems with non-cumulative topologies. These similarities indicate where non-cumulative systems might be investigated as systems with equivalent cumulative topologies. Note that a similar parametric study is presented by Konakli et al. (2015); the results presented here are qualitatively similar, and extend the parametric study to a larger system, as well as considering multiple component-level actions rather than a single system-level action.

The systems considered consist of $n = 5$ components, whose performance is determined by a set of $n_X = 5$ Gaussian random variables, with the following joint distribution:

$$\mathbf{f} \sim \mathcal{N}(\boldsymbol{\mu}_F, \boldsymbol{\Sigma}_F) \quad \boldsymbol{\mu}_F = \mathbf{1} \quad \boldsymbol{\Sigma}_F = [(1-\rho)\mathbf{I} + \rho\mathbf{1}\mathbf{1}^{\mathrm{T}}] \tag{5-29}$$

where $\mathbf{1}$ denotes a vector of ones of appropriate length (in this case, length five) and $\mathbf{I}$ denotes the identify matrix of an appropriate size (again five in this case). The correlation coefficient ρ

defines the mutual covariance of any pair of these random variables, and is left as a parameter to be varied in this study.

Measurements of this random field of the form of Equation 3-5 are considered. Up to five measurements are possible, each corresponding to one of the underlying random variables. Depending on the choice Y of the subset of these measurements to use, the conditional distribution for $\mathbf{y}$ is:

$$\mathrm{p}_{Y|\mathbf{f}}(\mathbf{y}|\mathbf{f}) = \mathcal{N}(\mathbf{\Omega}_Y\mathbf{f}, \sigma_\epsilon^2\mathbf{I}) \tag{5-30}$$

where $\sigma_\epsilon^2 = 10^{-10}$, i.e., a negligibly small measurement error is considered. Matrix $\mathbf{\Omega}_Y$ is defined such that the entry in the ith row and the jth column is one if the ith entry of $\mathbf{y}$ corresponds to a measurement of jth entry of $\mathbf{f}$, and zero otherwise, capturing a local measurement topology.

Limit-state variables are defined following Equation 3-11 for two different cases of functional topology. The first case, termed the *absolute* case, assumes that the state of each component depends directly on the value of its associated random variable. This case is defined by:

$$\mathbf{\Xi}_G = \mathbf{I} \tag{5-31}$$

The second case, termed the *relative* case, assumes that the state of each component is defined by the difference between the value of its associated random field variable and the average of these variables for all components. This case is defined by:

$$\mathbf{\Xi}_G = \mathbf{I} - \frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}} \tag{5-32}$$

In both cases, $\mathbf{g}_0 = 0$. Component states are defined from the limit-state variables as in Equation 3-14. While the absolute case corresponds with the more traditional reliability analysis framework (e.g., Ditlevsen and Madsen, 1996), the relative case captures certain situations of interest, for example that of differential settlements between foundations, where it is the difference between individual foundation settlements and the average structural settlement that is of interest, rather that the absolute settlement of any one foundation (e.g., Glisic et al., 2005).

The loss functions used depend on the system's management and network topologies. Following the prototypical examples of Section 4.2, binary actions are considered for each system component. For the cumulative system, the loss function of Equation 4-15 is used, with $C_{f,i} = 1$ for each component and $C_{r,i}$ selected such that $C_{r,i}/C_{f,i} = P_{f,i}$ for each component under the different cases for the limit-state variable definition. For the series and parallel systems, Equation 4-17 is used with effective state functions as in Equations 4-18 or 4-19 respectively. The system-level failure cost is set to $C_f = 5$, and the repair costs are the same as for the cumulative system.

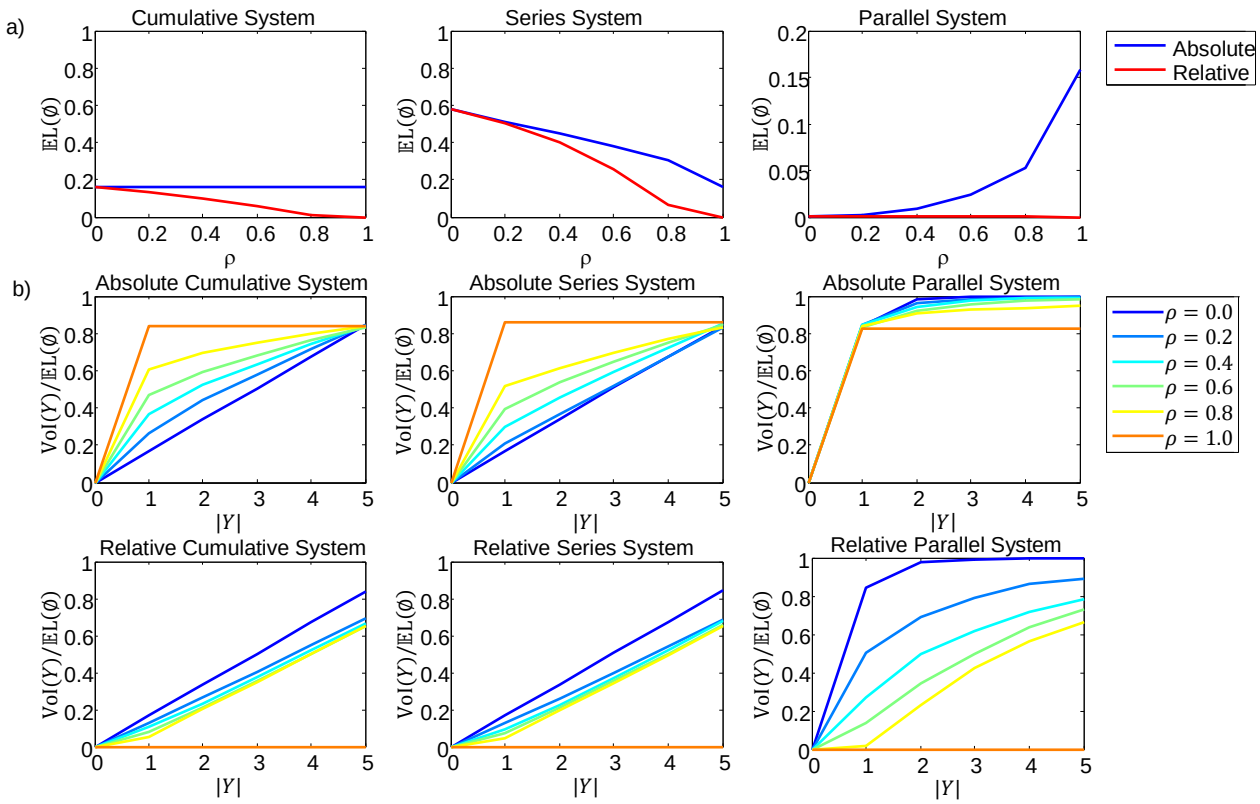

Figure 5-1: a) Prior expected loss as a function of correlation for various system topologies; b) VoI scaled by the prior expected loss as a function of number of measurements for various system topologies and correlation coefficients.

The results of this investigation are shown in Figure 5-1. First, in Figure 5-1a, the prior expected loss is evaluated for the three system topologies (cumulative, series, and parallel) and two limit-state definition cases (absolute and relative), under six settings for the correlation coefficient ρ (0, 0.2, 0.4, 0.6, 0.8, and 1). Note that for the absolute cumulative system, the prior expected loss is constant with respect to the correlation, since marginal component failure probabilities are not affected by ρ. In the relative cumulative system, however, this correlation affects the failure probability since higher correlation means that variables for individual components will be closer to the mean value of these variables for all components, and therefore failure probabilities are smaller under the relative failure definition. At $\rho = 0$, these two cases are equivalent. For series systems, the prior expected loss decreases as the correlation increases; the reliability of series systems is known to increase as their component behaviors become more highly correlated. For parallel systems, the opposite holds true in the absolute case, since as correlation increases, more components will fail together, making the parallel topology less reliable. In the relative case, however, prior expected cost remains low, since as correlation increases individual component failure, and therefore system failure, becomes less likely. Note that for all topologies, at $\rho = 1$, the system's components all operate or fail together, and the system behaves as a single

component, regardless of topology. The prior expected loss is thus the same across topologies for $\rho = 1$.

Figure 5-1b plots the VoI for measurement sets of different sizes, under each system topology, limit-state definition, and correlation level considered. The VoI is scaled by the prior expected loss for each system as shown in Figure 5-1a to allow for easier comparison. Note that, as the systems are defined, there is no distinction between components; they all have the same marginal failure probabilities and management costs, and the random variables affecting each are correlated with those of the others to the same degree. Thus, there is no difference between the VoI provided by any subset of measurements of the same size in this system. Therefore, optimization of sensor placements is unnecessary; only the number of sensors is relevant to VoI.

In the absolute cumulative system, for fully correlated components ($\rho = 1$), the maximum VoI is obtained from the first sensor placement; since all components perform exactly the same, measuring any one component gives full information about the system state. For uncorrelated components ($\rho = 0$), no information is shared between components, and so VoI grows linearly with the number of measured components. Between these extremes, the relative value of the first measurement compared to subsequent measures increases as the correlation increases. Results for the absolute series system are qualitatively similar, although it should be noted that the magnitudes of VoI are different due to the different prior expected losses.

For the absolute parallel system, for fully correlated components, the performance of the scaled VoI matches that for other topologies, since all components are in effect acting as one no matter the topology. However, in the case of lower correlation, the scaled VoI of multiple measurements is higher than in the fully correlated case; this is the only topology exhibiting this effect. This is due to the fact that, with multiple measured components, there is a greater chance of observing an operational component, and thus not having to take any costly management actions. This chance increases as the correlation between the components decreases. These results indicate that, regardless of correlation, a single measurement in a parallel system can have comparatively large relative benefits compared to single measurements in other system topologies. However, it should be noted that the magnitude of the VoI is generally lower in parallel systems due to the lower prior expected loss, as shown in Figure 5-1a.

For the relative limit-state definition case, note that there is never any expected loss or VoI when $\rho = 1$, since there is no possibility for component or system failures. Furthermore, for $\rho = 0$, the scaled VoI is the same under the absolute and relative limit-state definitions, since for independent random variables, these two definitions are equivalent. For other correlation levels, VoI is comparatively lower in the relative case than in the absolute case, and decreases with increasing correlation across all topologies. A likely reason for this is that the measurement information must be used to estimate both the local random field value and the mean of all random field values in order to determine the component state. VoI provided by sensors is therefore lower, as with few measurements the estimate of the mean may be poor.

Several general conclusions can be drawn from these results. Under the relative limit-state definition, the low VoI provided by single measurements may be a hindrance to greedy optimization; if the first measurements to be selected all have a relatively low or negligible VoI, there will be little to distinguish between them, and thus the first selected measurement may lead to the selection of an overall suboptimal set. On the other hand, in parallel systems under the absolute limit-state definition, the relative VoI of single measurements is very high, which might indicate that greedy optimization will perform well.

Furthermore, the qualitative similarities between the scaled VoI in the cumulative and series system topologies suggests that the latter might be approximately analyzed using an appropriate equivalent definition for the former. That is, by using the components of the series system to virtually construct an equivalent cumulative system, VoI can be more efficiently computed for measurements in the cumulative system, using the techniques of Section 5.2.2. Then, using the ratio of the prior expected losses in both systems (note that, although still computationally intensive to evaluate for a series system, this expected loss need only be computed once), the VoI in the cumulative system can be transformed into an approximate VoI for the series system. The development of approximate techniques for VoI evaluation in series systems based on this approach is left as a topic for future work; a preliminary investigation into the feasibility of such an approach is included in Section 11.2.

Finally, Figure 5-2 presents an empirical investigation of the computation time for VoI evaluation. For this analysis, the evaluation of the VoI for a single measurement on any component is considered as described in this section. The number of system components is varied while maintaining the system topology, and $\rho = 0.5$ is used across all system sizes and topologies. Computation is performed on a laptop computer with a 2.6 GHz Intel® Core™ i7-6700 CPU with 32 GB RAM. Code for evaluation is implemented in MATLAB R2016b. In Figure 5-2a, the computational time for each evaluation is plotted, with this time being the average across ten runs. For comparative purposes, an exhaustive search is also implemented, i.e., every combination of joint state and action is considered. Comparing to the dotted line, which represents an exponential growth rate of $\mathrm{O}(8^n)$, this exhaustive search seems to be taking exponentially greater time as the system size increases, although there is some variation in the rate of increase. Evaluation in parallel and series systems, although made easier through the methods discussed in Section 5.3.2, appears to also be increasing at an exponential rate as well. For cumulative systems, however, the computational time is actually decreasing, although this is likely an approximation error due to the very small computational times. In Figure 5-2b, the analysis of the cumulative topology is extended to larger systems, and the rate of computation time increase can be more clearly seen. Below about 100 components, the computation time is not very sensitive to the system size. However, above this number, the rate of growth increases, being approximately the same as that of the dotted line corresponding to an $\mathrm{O}(n^3)$ increase rate. Overall, the time to evaluate VoI in a cumulative system with five thousand components is roughly the same as the time needed in a non-cumulative system with five components. Of

course, these results are dependent upon the implementation of the code, and may not be valid beyond the range of system sizes considered.

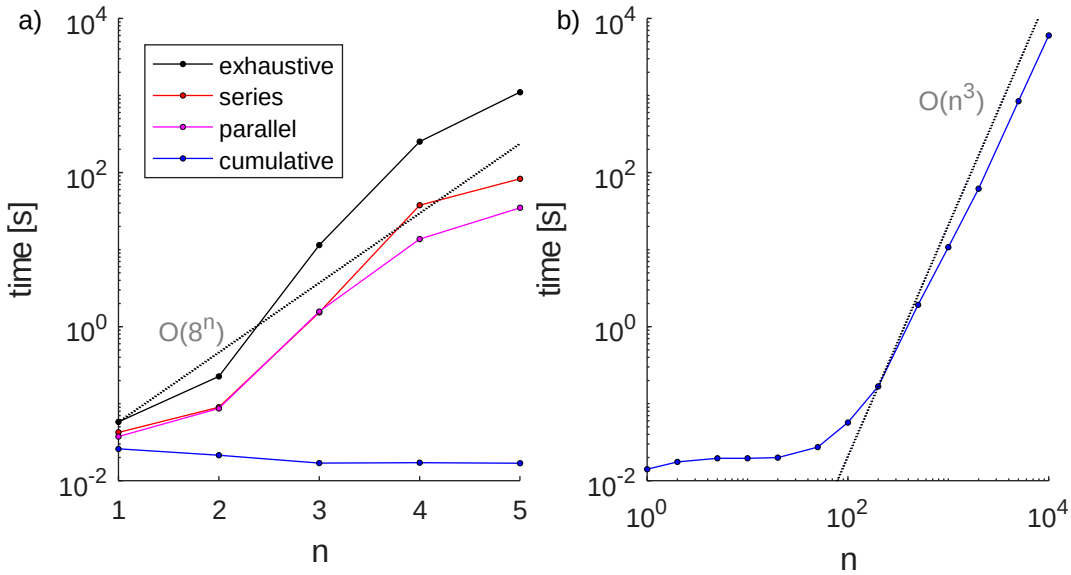

Figure 5-2: a) Computational time for VoI evaluation versus number of components for several system topologies, with dotted line indicating exponential growth rate; b) Extended plot of computation time versus number of components for cumulative systems, with dotted line indicating cubic growth rate.

Chapter 6

Comparative Examples for Sensor Placement

In this section, the application of the VoI metric is demonstrated in two example problems, based on the motivating example of Section 1.1. First, in Section 6.1, a one-dimensional random field problem, inspired by the loading on a structural element such as a beam, is examined. In this problem, various topologies for the system's functionality and management are considered for defining the problem's probabilistic model structure and loss function. Under these definitions, the VoI obtained by a single measurement at various locations on the one-dimensional domain is computed, and used to optimize the choice for where to locate this single measurement. Note that placing this first measurement is the first step in either greedy offline or online sensor placement, as discussed in Section 2.2.1. Additionally, results obtained for sensor placement using alternative metrics, such as the prediction error and misclassification rate, are presented and compared to results of the VoI metric under various topologies.

Next, in Section 6.2, a two-dimensional random field problem, motivated by a corrosion prevention application, is considered. Various system topologies are considered, and the VoI under each is used to optimize the location of sensors. In this example, multiple sensor placements are considered, to illustrate the application of the greedy offline optimization approach. Additionally, the use of the random field prediction error metric is demonstrated for this problem, and a comparison between online and offline sensor placement approaches is made. Overall, the goal of this chapter is to demonstrate the full procedure for the application of the VoI metric for optimal sensor placement on intuitively understandable examples. This serves both to illustrate the applicability of the approach, and to provide the reader with additional insight into how the VoI metric handles different problems, to facilitate problem-specific intuition.

6.1. One-Dimensional Example in Structural Loading

The first example application of the VoI metric to optimal sensor placement considers a one-dimensional structure of length $D = 10$ meters that must support an uncertain distributed load, e.g., a cross-section of a roof under snow loading or of a geotechnical structure under surface load. The loading is modeled as a Gaussian random field, i.e., a Gaussian process model in one-dimension is used to describe the loading $\mathrm{f}(x)$ at location x along the structure, quantified as a

linear pressure in units of kN/m. The Gaussian process model is constructed as in Equation 3-2 with a linearly increasing mean function:

$$\mu(x) = 15\,\frac{\mathrm{kN}}{\mathrm{m}} + \left(0.5\,\frac{\mathrm{kN}}{\mathrm{m}^2}\right)x \tag{6-1}$$

and with a covariance function as in Equation 3-3, with a standard deviation function describing a linear decrease:

$$\sigma(x) = 2.5\,\frac{\mathrm{kN}}{\mathrm{m}} - \left(0.2\,\frac{\mathrm{kN}}{\mathrm{m}^2}\right)x \tag{6-2}$$

and with a correlation function that uses the common squared exponential correlation form (Rasmussen and Williams, 2006):

$$\rho(\mathbf{x}_i, \mathbf{x}_j) = \exp\left[-\frac{1}{2}(\mathbf{x}_i - \mathbf{x}_j)^{\mathrm{T}}\boldsymbol{\Lambda}(\mathbf{x}_i - \mathbf{x}_j)\right] \tag{6-3}$$

where, in this case, the coordinate vector $\mathbf{x}$ has a single element. This correlation structure describes an asymptotic decrease of correlation, i.e., movement towards the independence of random variables as the distance between these variables increases. The rate of this decrease for different dimensions is controlled by the scaling parameter matrix $\boldsymbol{\Lambda}$. In this problem, $\boldsymbol{\Lambda} = \mathbf{I}\lambda^{-2}$, with correlation length scale parameter $\lambda = 1/\sqrt{2}$ m.

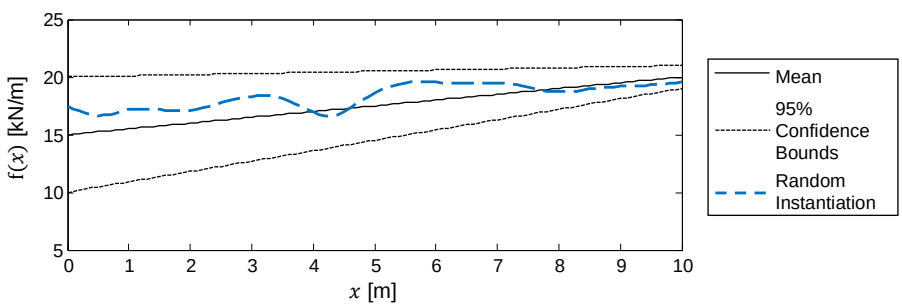

Figure 6-1: Prior one-dimensional Gaussian process loading model, with an example instantiation of the loading field.

The prior model of the random loading field is depicted in Figure 6-1. The prior mean and 95% confidence boundaries illustrate the effects of Equations 6-1 and 6-2. Note that, in an expected sense, the structure is more heavily loaded towards the right, but this loading is more uncertain towards the left. The effect of the correlation structure of Equation 6-3 is illustrated by the random instantiation of the field shown in the figure; the smoothness of the field is determined by the correlation structure and correlation length, with a longer correlation length leading to a smoother field.

The domain is discretized into $n_X = 100$ segments, each of length $\Delta x = 0.1\text{m}$, to define the set of discretized locations X, and thus the random loading field vector $\mathbf{f}$. A loading gauge is to be placed at one of these discretized locations, with position x_y along the structure, and will take a noisy measurement of the local loading. This measurement is described as in Equation 3-5 with $\mathbf{\Omega}_Y$ being a row vector with n_X entries, all being zero except for a one in the position corresponding to the location of x_y. The measurement noise is a single random variable with zero mean and standard deviation $\sigma_\epsilon = 0.1\text{kN/m}$.

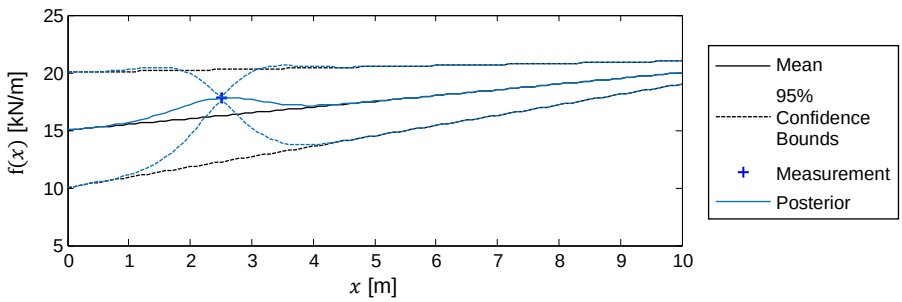

Figure 6-2: Posterior one-dimensional Gaussian process loading model, conditional to an example observation.

An example of the updated posterior model for $\text{f}(x)$ conditioned on an observation by this gauge at location $x_y = 2.5\text{m}$ is given in Figure 6-2. Due to the correlation structure, a measurement at one point updates the distribution not just at this point, but at nearby locations as well; this results in the posterior confidence bounds for this random field being narrower in the vicinity of the measurement.

The following subsections examine different management problems for the structure resulting from this random loading field. The characteristics of these problems determine the limit-state variable definitions used, as well as the possible management actions and resulting loss functions for the VoI metric. Across these problems, three methods for formulating the loss function are considered, corresponding to special cases of the VoI. The first, which is based on *estimation* of the underlying random loading field (or a transformation of this field) follows from the prediction error metric as discussed in Section 4.4.1. The second, based on *classification* of the binary component or system states, follows from the misclassification rate metric as discussed in Section 4.4.2. The third, based on the *maintenance* of the system, follows from the prototypical management problems discussed in Section 4.2. Specific forms of these loss functions depend on the problem being investigated, and are described in more detail below.

Finally, note that a related problem is investigated by Bednarz et al. (2013a, 2013b), where strain gauges installed on a beam are used, along with known relationships between loading, bending moment, and strain, to determine the location and magnitude of a point load applied to the beam.

In that case, no prior loading model is used, and there is no consideration for optimizing strain gauge locations. However, many of the same static analysis principles apply to both problems.

6.1.1. Cumulative System

It is first assumed that the system has a cumulative topology. Under this assumption, if the loading at any point along the structure exceeds threshold $f_{\max}$, a failure occurs at that point. For example, in a geotechnical structure, collapses may occur if the loading exceeds the soil's bearing capacity locally. The limit-state variables are defined as in Equation 3-11 with $\mathbf{\Xi}_G = -\mathbf{I}$ and $\mathbf{g}_0 = \mathbf{1} f_{\max}$, i.e., these represent the difference between the threshold and the local loading. Local states are then defined for each location following Equation 3-14. For *estimation* of the loading field, the loss function of Equation 4-31 is used. For *classification* of the local states, that of Equation 4-37 is used, with a factor of $1/n$ such that the average misclassification across local states is represented. For *maintenance* of the system, the prototypical cumulative system loss function of Equation 4-15 is used, with local binary actions corresponding to the choice to reinforce the structure at each discrete point to prevent failure. For this problem, $C_{f,i} = c_f \Delta x$, and $C_{r,i} = c_r \Delta x \; \forall i$, where $c_f = \$100/\text{m}$ and $c_r = \$10/\text{m}$ are the cost densities for failure and reinforcement respectively.

Figure 6-3 presents the results of VoI evaluation in this system, plotting the VoI of a measurement versus its position x_y in the structure, under the three loss function definitions considered. Since, under each loss function, the VoI is measured in a different unit, each VoI is normalized by its maximum value under each loss function, to allow comparison of the dimensionless scaled VoI across different loss function. These maximum values, corresponding to the VoI achieved by the optimal placement in each case, are listed in Table 6-A for reference for this and the other problems considered in Sections 6.1.1 through 6.1.3 (Section 6.1.4 is not included, since the system management problem considered is somewhat different from the other subsections). Also note that Table 6-A lists threshold values $f_{\max}$ which are different for these cases; this is done in order to ensure an interesting management problem in each case, e.g., if the same threshold were used in a parallel system as for any of the other systems, the probability of system failure would be negligible.

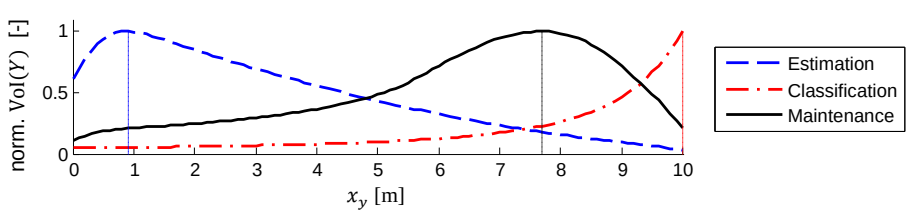

Figure 6-3: VoI metric results for a cumulative system topology.

For the objective based on field *estimation*, the optimal sensor placement location is towards the left of the domain, since the prior variance of the loading is highest in this area. Note that, rather than being placed to the extreme left side, the sensor is placed more centrally in order to take advantage of local correlations to reduce posterior uncertainty to both the left and right of the sensor location. For the objective associated with state *classification*, the placement is to the extreme right, since the prior misclassification rate for individual segments is close to its maximum of 50% here. A measure at this location will therefore best support appropriate classification for these components, and therefore reduce the overall misclassification rate the most. For the objective supporting system *maintenance*, relatively high probabilities of local failures to the right of the domain lead to the optimal action being to pre-emptively reinforce these sections. To the left of the domain, because of the lower probability of failure, it is better to do nothing. In between these extremes, additional information is most helpful in deciding which actions to select, and therefore VoI is higher.

Table 6-A: VoI metric results for the systems of Sections 6.1.1 through 6.1.3.

$\mathrm{VoI}(Y^*) =$	Threshold ($f_{\max}$) [kN/m]	Estimation (sum of marginal variances)	Classification (misclassification rate)	Maintenance (expected loss reduction)
Cumulative	20	65.14	0.4344	10.26
Global	18	65.14	0.0236	16.51
Series	21	65.14	0.0251	11.93
Parallel	14	65.14	0.0977	34.22

6.1.2. Global System

Next, a global system is examined, where the system as a whole will fail if the total loading on the structure exceeds a specified threshold, calculated as Df_{max}. Reinforcement of the entire structure is required to prevent collapse. To investigate this case, it is interpreted as a trivial example of a system with a cumulative topology having only $n = 1$ component. A single limit-state variable for the system is defined as in Equation 3-11 with $\boldsymbol{\Xi}_G = -\mathbf{1}^{\mathrm{T}}\Delta x$ and $\mathbf{g}_0 = Df_{\max}$, representing the difference between the total loading on the structure and the threshold on this total load. A single state variable is defined following Equation 3-14. For *estimation* of the load, the loss function of Equation 4-31 is again used. For *classification* of the global state, the loss of Equation 4-37 is again used, this time for the single state variable. For *maintenance* of the system, the prototypical loss function of Equation 4-15 is again considered for this single component, with the cost for a failure now being $C_f = c_f D$ and the cost for a reinforcement action on the entire system being $C_r = c_r D$.

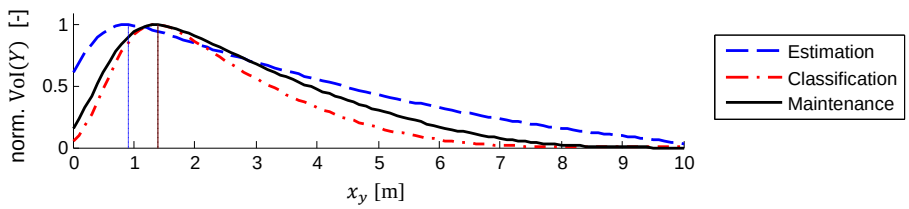

Figure 6-4: VoI metric results for a global system topology.

Figure 6-4 presents sensor placement results for this system. The placement based on *estimation* is unchanged. For the objectives based on *classification* and *maintenance*, the optimal measurement locations are the same, and are near that for the random field *estimation*. This can be understood by noting that, for a single component, accurately predicting its state is an essential step to choosing the appropriate posterior management action. Furthermore, since the state is a function of the total loading, a more accurate prediction of this total will allow for a more accurate prediction of the state, and a more accurate prediction of the loading at each individual location will lead to a more accurate prediction of the total. Therefore, in this case, the objective for random field prediction provides a good proxy for those of state classification and optimal management for this single component. This represents a case where a more computationally efficient loss function can be substituted for a more complicated function, and still give nearly the same results in terms of the optimal placement location.

6.1.3. Series and Parallel Systems

Now, series and parallel system topologies are considered, as illustrative examples for non-cumulative topologies. Specifically, while local failures can occur as in Section 6.1.1, these will, in the series case, lead to a failure of the entire system, while in the parallel case, all discrete sections of the structure must locally fail for the system to fail. The series case approximates one of progressive collapse, where any local structural failure can lead to failure of the entire system; for ease of analysis, the exact method of this collapse is ignored, and only its initiation as a single local failure is modeled. While the parallel case is not commonly encountered in structural systems, it is nevertheless presented for the sake of completeness. Such a topology may occur in a non-structural system, e.g., a communication system can be considered operational so long as communication is possible by at least one channel.

Local limit-state variables are defined, as in Section 6.1.1, following Equation 3-11 with $\boldsymbol{\Xi}_G = -\mathbf{I}$ and $\mathbf{g}_0 = \mathbf{1} f_{\max}$, and local states are again defined following Equation 3-14. For *estimation*, the values of the local loading are again considered the quantities to be predicted, following the loss function of Equation 4-31. For classification, the binary state of the system is to be predicted, following the loss function of Equation 4-38 with the system state defined by Equation 3-19 for the series system and Equation 3-20 for the parallel system. For *maintenance* in the series system, the loss function of Equation 4-17 is considered, with the effective state of

Equation 4-18 and with failure cost equal to the global failure cost of $C_f = c_f D$ and repair costs for each component equal to the local repair costs of $C_{r,i} = c_r \Delta x \; \forall i$. For *maintenance* in the parallel system, because a single action has the possibility to ensure the operation of the entire structure, a single reinforcement action is considered, following the loss function of Equation 4-16, with the system state of Equation 3-20, and with the same failure cost as the series system and a reinforcement cost of $C_r = c_r D$.

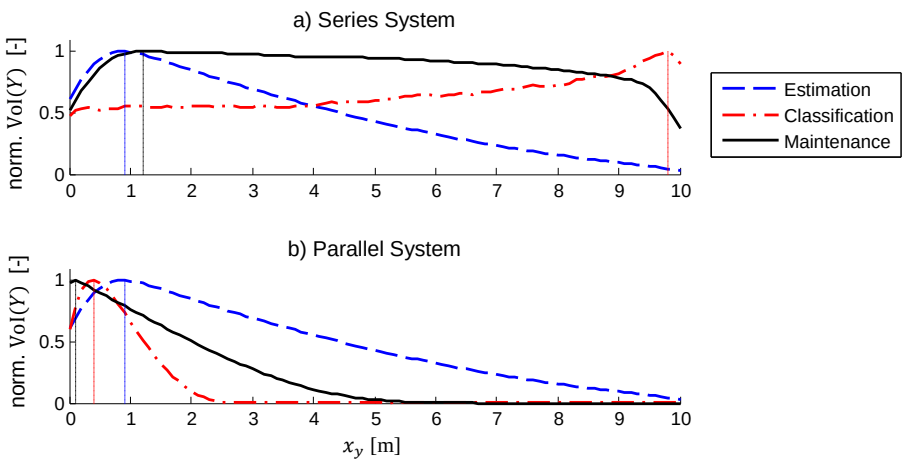

Figure 6-5: VoI metric results for a) series and b) parallel system topologies.

Figure 6-5 presents sensor placement results for these systems. Note as before that the optimal placement for random field *estimation* remains unchanged. The placement for state *classification* is either to the far left or right, depending on the topology. In the series system, this placement is to the extreme right; intuitively, observing a single failure (i.e., observing a load value above the local threshold) in a series system will lead to the conclusion that the system has failed, and the most likely place to observe such a failure is towards the right. Conversely, in the parallel system observing a single operational component ensures that the system is operational, and the most likely place to observe such a component is towards the left side. The same reasoning does not hold true for the system *maintenance*, however, as in the series system, the a-priori optimal action is to perform the reinforcement activity for most of the structure to prevent a costly system failure. To the left side, where local failures are unlikely, an additional measurement provides valuable information as to whether or not reinforcement is needed here. In the parallel system, a similar optimal sensing position is chosen, for a different reason; here, the loading must exceed the threshold everywhere to cause a failure. If this loading is observed to exceed the threshold towards the left side, then it is likely it will exceed the threshold everywhere else (since the left has the lowest mean and highest variance), and so a reinforcement action is probably needed. Otherwise, if a low loading value is observed, no action is necessary. Finally, note that although the objective based on *estimation* provides a reasonable proxy to that based on *maintenance* in both of these systems, this is due to a complicated interaction between the structure of the random field and the system topology, which cannot be expected to hold in general.

6.1.4. Multiple Limit-State Functions

In this last one-dimensional example, consider that the structural element in question is a simply supported beam. The performance of this beam is governed by three constraints:

1. The total loading on the beam must not exceed 200kN;
2. The loading on either support must not exceed 100kN;
3. The moment at the center of the beam must not exceed 250kNm.

A system failure cost of $C_f = \$1\text{M}$ will be incurred if any of these limits are exceeded. For each of these conditions, an associated intervention action can be applied to locally strengthen the beam or its supports and prevent failure; these are associated with cost $C_{r,1} = \$5\text{k}$ to prevent failure due to the total load, costs $C_{r,2} = C_{r,3} = \$2.5\text{k}$ to prevent failure due to loadings at the left and right supports, and cost $C_{r,4} = \$7.5\text{k}$ to prevent failure due to bending in the center.

This problem statement defines a management decision-making problem for a four-component system in series. The components of this system are virtual, corresponding to four different failure scenarios for the beam. The occurrence of these scenarios is indicated using four limit-state quantities, associated with the total beam loading, the left support reaction, the right support reaction, and the moment at the center. The limit-state variables are defined following Equation 3-11, with each row of matrix $\boldsymbol{\Xi}_G$ encoding the (negative) influence line that translates the structural loading field to either the total load, left or right reactions, or moment at the center. These influence lines are plotted in Figure 6-6, following basic principles of static analysis (e.g., Hibbeler, 2009). Threshold values for these engineering quantities of interest are encoded as $\mathbf{g}_0 = \{200\text{kN}, 100\text{kN}, 100\text{kN}, 250\text{kNm}\}$. Local states are defined following Equation 3-14, and the effective system state considering the maintenance actions is expressed as in Equation 4-18.

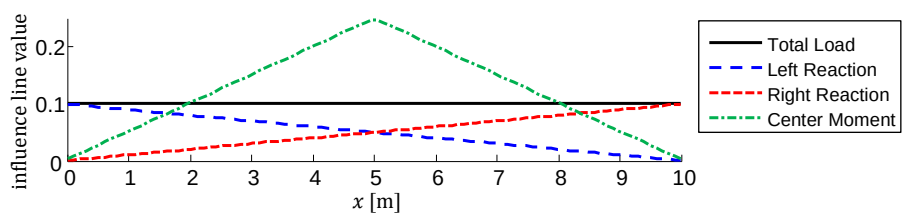

Figure 6-6: Influence lines for structural loading.

From this problem, VoI is calculated under two different cases, with two different loss functions in each case. First, each of the failure conditions is considered separately. This is analogous to managing the system considering each failure mode on its own, and determining an optimal sensor placement for guiding the system's management under each loss function (i.e., modeling the system as a cumulative rather than a series system). Optimal placements are then determined by either considering the loss function of limit-state value estimation for each of the engineering

quantities of interest (total load, left reaction, right reaction, and moment at the center), following Equation 4-33 for a single limit-state variable, or considering the management of a single-component system based on this quantity, following Equation 4-15 for a single component. Second, the failure conditions are considered together, and the VoI and optimal placement are determined following the loss function of Equation 4-33 using the vector of four limit-state variables and the loss function of Equation 4-17, considering a series system of the failure conditions.

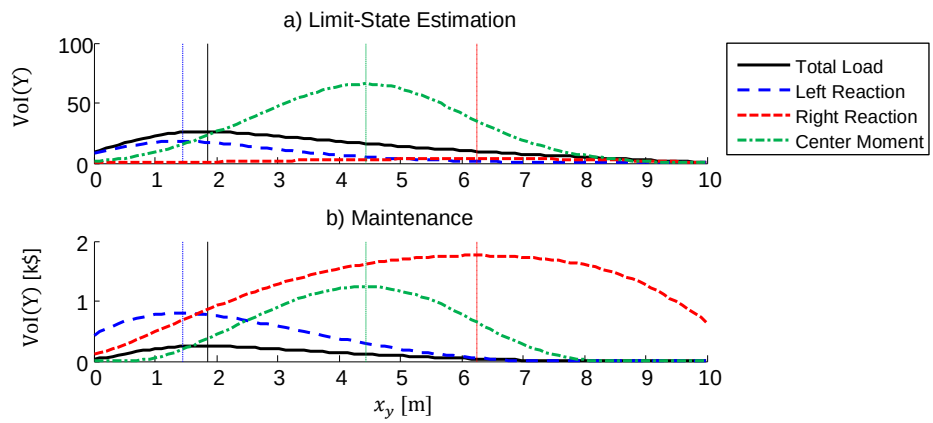

Figure 6-7: VoI metric results considering failure conditions separately, for a) estimating limit-state variables separately and b) managing the system under each failure condition.

Figure 6-7 shows the results for VoI and for the resulting sensor placements, indicated by vertical lines at the maximum value of VoI for each engineering quantity considered. Figure 6-7a shows the VoI for prediction of the limit-state quantities, i.e., for the prediction of the total load, left reaction force, right reaction force, or moment at the center (since the thresholds on these quantities are known). The effects of the influence lines can be seen here, i.e., the optimal measurement for determining the right reaction force is to the right, the optimal for the left reaction force is to the left, and the optimal for the center moment is near the center. Note that these positions are all somewhat skewed towards the left side, rather than symmetrical about the center, and that the placement position for measuring the total load is near the left as well. This is due to the effect of the prior uncertainty in the loading field, which is higher towards the left.

Figure 6-7b shows the VoI of sensor placements for managing the system considering only one of the failure conditions at a time. Note that, although the VoI values are different, the optimal sensor placement locations are the same as for the limit-state variable prediction in Figure 6-7a. This can be understood by noting that the state of each virtual component, and thus the choice of optimal action to manage that component, is a function of the limit-state variable associated with this component. Thus, a better prediction of this limit-state variable will better predict the state, allowing for better choices of posterior management actions, and therefore a higher VoI.

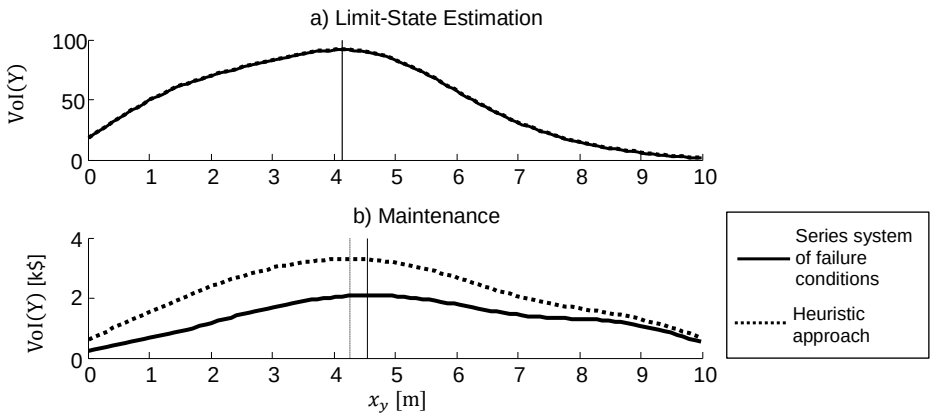

Figure 6-8: VoI metric results considering failure conditions together, for a) estimating the limit-state variables together and b) managing the series system of failure conditions.

Figure 6-8 shows the VoI and the resulting sensor placements when all the failure conditions for the system are considered together, i.e., as a series system of the virtual components associated with each failure condition. Figure 6-8a presents results for prediction of the limit-state variables following Equation 4-33. Note that the optimal placement location for reducing prediction error across all limit-state variables is different from that for reducing errors for any individual limit-state variable, as in Figure 6-7a. However, there is a relationship between these goals; the curve in Figure 6-8a represents using the sum of the curves for prediction VoI of the individual limit-state variables in Figure 6-7a. This relationship is the result of the decomposability of VoI across components in cumulative systems illustrated in Equation 5-5. In this case, the system management curve exactly matches the sum of component management curves because the trace of the covariance matrix for G is equal to the sum of the variances of each element of G.

Figure 6-8b shows the VoI and associated optimal sensor placement location for managing the system considering all failure conditions together, i.e., as a series system of the virtual components associated with each failure condition. In this case, the optimal sensor placement does not match that of Figure 6-8a, as was the case for the individual limit-states considered separately. Moreover, a heuristic approach of adding the individual VoI curves of Figure 6-7b (effectively considering the system topology to be cumulative rather than series) does not equal the VoI curve for the series system, and the resulting sensor placements also do not agree. Nonetheless, in this case the heuristic represents a reasonable approximation; the VoI provided by the sensor placement indicated using this heuristic is 98% of the VoI of the true optimal sensor placement location. This provides some anecdotal evidence supporting the approximate approach discussed at the end of Section 5.4, as the true optimal sensor placement in this series system is reasonably approximated by the placement resulting from summing the VoI of the measurements for each of component, i.e., by treating it as a cumulative system and using Equation 5-5.

6.2. Two-Dimensional Example in Corrosion Prevention

The second example application of VoI to sensor placement is motivated by the study of an historic steel mill in Pittsburgh. In this problem, a 20m by 20m section of steel ceiling is subjected to corrosion due to water vapor produced as a by-product by machinery operating underneath it. For the purpose of this example, corrosion is induced by a sustained high relative humidity (RH) under a certain section of the ceiling. Corrosion will eventually lead to ceiling leakage, causing damage to the building interior and its contents. The building manager has the option to measure long-term average RH using sensors, as well as coat the ceiling to prevent corrosion and resulting damage.

Gaussian process models in two-dimensions are used to model the random field of the long-term average RH beneath the ceiling and the capacity of the steel ceiling to resist corrosion. The RH field, also termed the demand field or F_d, is modelled following Equation 3-2, with a baseline mean of 40% RH and a standard deviation of 5% RH throughout the building. An isotropic square exponential correlation function, as in Equation 6-3 with correlation length scale parameter $\lambda = 4$m, defines the correlation structure. Furthermore, in addition to the constant baseline humidity, internal equipment produces localized plumes of higher humidity. In the center of the lower-left quadrant of the building (as viewed from above), equipment produces a plume of increased humidity, with average peak of 20% RH above baseline, with standard deviation of 5% RH. In the center of the upper-right quadrant, equipment produces a plume with a lower average peak (10% RH) but a higher standard deviation (15% RH), corresponding to a less certain humidity output from this equipment. The overall demand field $\mathrm{f}_d(\mathbf{x})$ is the sum of the baseline field and the fields due to plumes from equipment. Figure 6-9a plots the prior mean $\mu_d(\mathbf{x})$ of the demand field, and Figure 6-9b plots its standard deviation $\sigma_d(\mathbf{x})$.

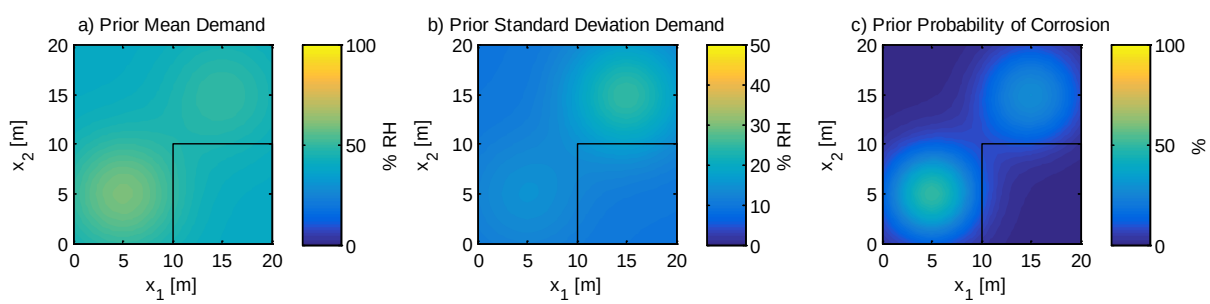

Figure 6-9: a) Prior mean of the random demand field; b) Prior standard deviation of the random demand field; c) Prior probability of corrosion for each location.

A simple model for corrosion initiation is used, with corrosion occurring if the long-term average RH below a section of ceiling exceeds to the local capacity of the steel to resist corrosion. This capacity field F_c is modeled as a Gaussian process, with a mean of 60% RH throughout, based on the critical relative humidity for corrosion initiation in steel (Badea et al., 2011). A standard

deviation of 5% RH throughout is assumed to account for other variables, such as ambient pollution or material defects, which influence the initiation of corrosion. An isotropic square exponential correlation function with $\lambda = 1\text{m}$ defines the correlation structure for this field.

For analysis, the ceiling is discretized on a 0.5m square grid, defining X with $n_X = 3200$ (i.e., 1600 spatial coordinates each for 2 index coordinates distinguishing the demand and capacity fields). The vectors for the random fields are concatenated to define the random field variable vector, following Equation 3-4, as:

$$\mathbf{f} = \begin{bmatrix} \mathbf{f}_c \\ \mathbf{f}_d \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \boldsymbol{\mu}_{F_c} \\ \boldsymbol{\mu}_{F_d} \end{bmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_{F_c} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}_{F_d} \end{bmatrix} \right) \tag{6-4}$$

where $\boldsymbol{\mu}_{F_c}$ and $\boldsymbol{\mu}_{F_d}$ are the mean vectors and $\boldsymbol{\Sigma}_{F_c}$ and $\boldsymbol{\Sigma}_{F_d}$ the covariance matrices for the capacity and demand fields, respectively, defined using the Gaussian process models for these fields. Note that the capacity and demand variables are modeled as being independent. Local limit-states for the ceiling are defined following Equation 3-11 with $\boldsymbol{\Xi}_G = [\mathbf{I} \quad -\mathbf{I}]$ and $\mathbf{g}_0 = \mathbf{0}$, i.e., the limit-state variable at each location is defined as the local capacity variable minus the local demand variable. Local states are then defined as in Equation 3-14, with each discretized section of the ceiling representing a component of the ceiling system (defining $n = 1600$ components), and its state indicating whether or not corrosion will initiate (i.e., $s_i = 0$ corresponds to corrosion initiation). Note that, since the demand field represents the long-term average RH within the structure, corrosion is assumed to initiate after a certain length of time under exposure to an average RH level that exceeds the local steel RH capacity. The prior local probabilities of corrosion initiation, i.e., the values of P_f for each of the roof sections, are shown in Figure 6-9c. Note that due to both the high average humidity in the lower-left quadrant and the high variability of humidity in the upper-right, both areas have relatively high prior probabilities of corrosion. Random realizations of the demand and capacity fields, as well as the resulting map for corrosion, are shown in Figure 6-10.

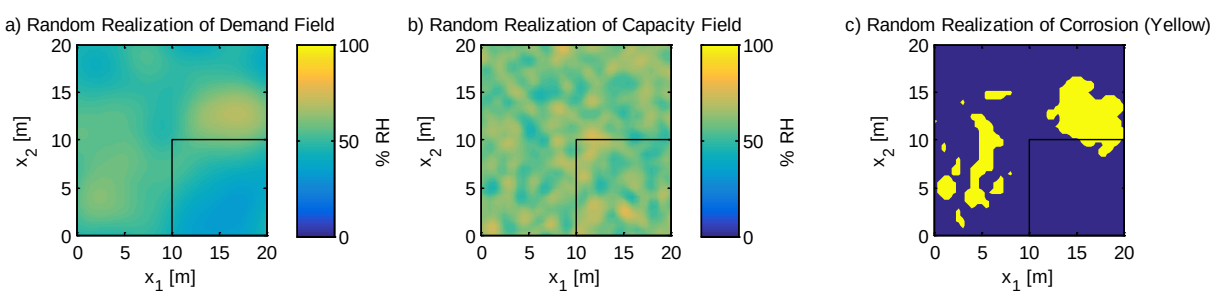

Figure 6-10: a) Realization of the demand field; b) Realization of the capacity field; c) Realization of the corrosion field (i.e., the component states), with yellow denoting corrosion.

Corrosion of the ceiling will incur some cost to the manager due to damage within the structure resulting from ceiling leakage. This cost is estimated at $\$100/\text{m}^2$ for most of the building, but in

the lower-right quadrant (outlined by a black square in Figure 6-9 and Figure 6-10), due for example to the location of sensitive electronics here, there is a higher consequence of $\$500/\mathrm{m}^2$. To avoid these consequences, the manager has the option to apply a coating to the ceiling, at a cost of $\$50/\mathrm{m}^2$, to prevent corrosion.

To support this decision, the owner has five humidity sensors to place in the building. Measurements of the sensors are modeled as in Equation 3-5, with the rows of $\mathbf{\Omega}_Y$ consisting of zeroes, except for a one in one of the latter 1600 locations, corresponding to a local measurement of the demand (i.e., the long-term average RH) at that location. Noise $\boldsymbol{\epsilon}$ is modeled as zero-mean independent noise with a standard deviation of 1% RH. Sensing cost is modeled as $\mathrm{C}(Y) = \mathbb{I}(|Y| > 5)$ with budget $b = 0.5$, corresponding to a strict limit of 5 sensors, with no cost per sensor (except in the case of Section 6.2.2, where linear cost functions are used). Data collected by these sensors over a period of time are used to update the long-term average humidity field model and identify optimal posterior management actions. It is assumed that there is sufficient time between when data are collected and when corrosion initiates for these actions to be taken to prevent the corrosion.

The following subsections examine five different cases for the network or management topology, sensor placement objective, and optimization approach. Sections 6.2.1, 6.2.2, and 6.2.3 consider a cumulative system, a series system, and a system with a single binary decision problem. Sections 6.2.4 and 6.2.5 revisit the cumulative system, with the former employing an alternative objective based on random field estimation, and the latter considering online rather than offline sensor placement optimization.

6.2.1. Cumulative System

First, consider a case where humidity sensors are placed offline to address the problem of optimal management of the ceiling when coatings are applied locally, i.e., actions are taken for individual sections of the ceiling. Using the prior probability of corrosion, shown in Figure 6-9c, the prior expected loss for the ith section $\mathbb{E}\mathrm{L}_i(\emptyset)$ is computed for all $i \in \{1, \dots, n\}$, as shown in Figure 6-11a. Note that in the lower-right quadrant, where the cost for corrosion is highest, the expected loss is correspondingly higher. Also note that prior expected loss is capped at $\$50/\mathrm{m}^2$, since this is the cost to coat the ceiling to prevent corrosion; for areas where the prior expected loss under the “do nothing” action is higher than this threshold, the prior optimal choice is to coat this area. The prior optimal action $a_i^*(\emptyset)$ is shown in Figure 6-11b for all $i \in \{1, \dots, n\}$, with yellow indicating that an area should be coated, and blue indicating areas where doing nothing is the prior optimal maintenance action.

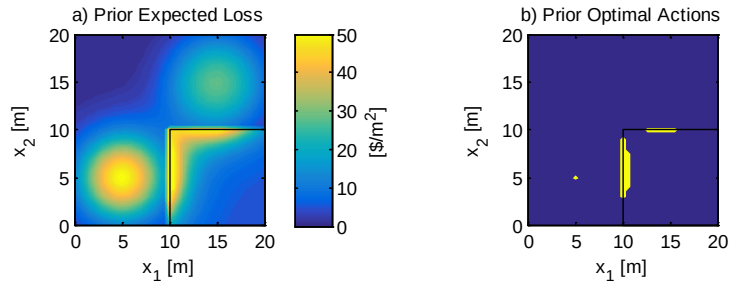

Figure 6-11: a) Prior expected loss for the management of the ceiling as a cumulative system; b) Prior optimal actions for the management of the ceiling, with yellow corresponding to a choice to apply coatings locally to prevent corrosion.

In this case, with local costs and maintenance actions, the loss function is that of a cumulative system, as in Equation 4-15, with appropriate failure (corrosion) and repair (coating) costs for each component (i.e., each ceiling section). Sensor placements are optimized via the VoI metric using this loss function and the offline greedy optimization algorithm, as discussed in Section 2.2.1.1. The results are shown in Figure 6-12, indicating the locations of sensors in the domain, as well as the order in which these sensors are greedily selected (indicated by numbers adjacent to the sensor locations) and the VoI provided by these sensors, using the order in which they were selected (e.g., the VoI for two sensors in the right plot indicates the VoI provided by two sensors at positions labeled 1 and 2 in the left plot).

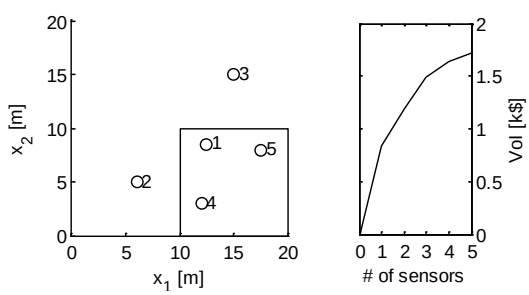

Figure 6-12: Placements for five humidity sensors optimized offline by the VoI metric for a cumulative system (left) and the VoI provided by these sensors (right).

Three of the five sensor measurements, including the first that was greedily selected, are located in the lower-right quadrant, for which the cost of corrosion is highest. These sensors are therefore best able to support the choice of posterior optimal actions to reduce the losses incurred in this quadrant. Additionally, the two sensors placed outside this quadrant are located near the centers of the humidity plumes produced by equipment within the structure. These measurements reduce uncertainty in the intensities of these plumes. The VoI provided by these five sensors is $1720, indicating that the expected system management cost is reduced by this amount due to data provided by these sensors supporting posterior optimal action selections.

6.2.2. Series System

To demonstrate sensor placement in a series system, consider a case in which any corrosion occurring in the lower-right quadrant of the structure will incur a cost of $50k, regardless of the extent of corrosion. This reflects a "zero tolerance policy" for corrosion in this quadrant, if for example it is assumed that any leaks in this quadrant will damage all the equipment therein. The consequences of corrosion anywhere outside this quadrant are negligible. Therefore, the lower-right quadrant exhibits a series system topology, where corrosion on any segment (or component of the system) results in failure of the entire system. The loss function of Equation 4-17 is used with the effective state as in Equation 4-18 and appropriate cost factors. The approximate methods outlined in Section 5.3 are used to evaluate VoI, which is optimized offline.

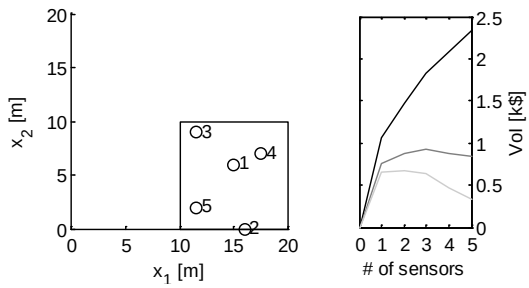

Figure 6-13: Placements for five humidity sensors optimized offline by the VoI metric in a series system of the lower-right quadrant (left) and the net VoI provided by these sensors using three alternative cost functions (right).

Figure 6-13 gives the optimal sensor placement by the VoI metric in this series system. Note that all sensors are placed in the lower-right quadrant, as in this case there is no reason to monitor humidity outside of the area that constitutes the system. While the prior optimal action is to coat the entire quadrant against possible corrosion, information gathered by the sensors can reduce the predicted probability of corrosion for certain sections, allowing these to safely remain uncoated, resulting in cost savings. The VoI provided by all sensors is $2330.

In this example, the VoI metric value or the net benefit of Equation 4-12 is also plotted to the right of Figure 6-13, for two alternatives sensor cost functions. At a cost of $300 per sensor, i.e., $C(Y) = \$300|Y|$ with $b = \$1500$, the optimal set, which maximizes net benefit (shown as the dark gray line), includes the first three sensors selected by the greedy algorithm, while for a higher cost of $400 per sensor (light gray line), only the first two sensors should be placed to provide the highest net benefit. This illustrates how in general VoI can be traded off against sensing cost to determine the optimal number of measurements to use to support decision-making.

6.2.3. Global Action

Rather than taking decisions of whether or not to coat individual sections of the ceiling, in this subsection the choice must be made whether or not to coat the entire ceiling. This would correspond to a situation where it is impractical to coat only certain sections of the ceiling, or where there is a flat fee for coating. The cost to coat the ceiling is $C_r = \$10\text{k}$ in this case, as opposed to $\sum_{i=1}^{n} C_{r,i} = \$20\text{k}$ for coating all sections individually. The loss function for this management problem is:

$$L(\mathbf{s}, a) = (1 - a)\left[\sum_{i=1}^{n} C_{f,i}(1 - s_i)\right] + aC_r \tag{6-5}$$

Note that the loss is no longer in a decomposable form. The management topology of the system, i.e., the single action choice applying to all components, couples decision-making, preventing decomposition. Fortunately, some of the efficient techniques of Section 5.2 for can still be applied to this problem. In particular, the conditional expected loss under the choice of $a = 0$ is decomposable, while under the choice of $a = 1$ the loss is fixed with respect to S. Therefore, an exhaustive search through the space of joint component states can be avoided. The approximate approach of Equation 5-25 is then used to take the expectation over observations.

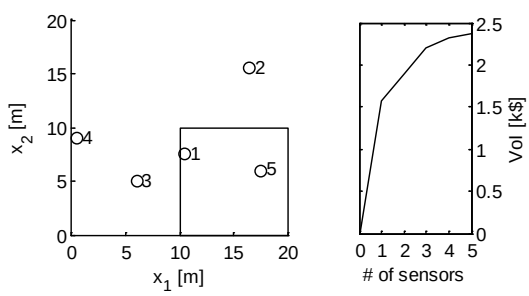

Figure 6-14: Placements for five humidity sensors optimized offline by the VoI metric considering a global management action (left) and the VoI provided by these sensors (right).

Using the VoI metric and this loss function, sensor placements are optimized offline, with the results shown in Figure 6-14. In this case, the prior optimal action is to do nothing, incurring a prior expected cost of $6600; collected data will prompt the manager to choose to coat the entire ceiling to prevent corrosion if higher-than-expected humidity is observed. The optimal sensor placement scheme includes a sensor in the upper-right quadrant, to monitor the region with the highest prior uncertainty, and two in the lower-right quadrant, where the consequences of failure are highest. Finally, it should be noted that similar arrangements of sensors that were tested in this case provided similar VoI, indicating that VoI is not as sensitive to the sensor locations in the case of global management decision-making as it is in the other cases investigated here. This is likely because the global decision is related to a weighted average of the humidity across the domain.

6.2.4. Comparison to Random Field Estimation

As an alternative to optimizing sensor placements using VoI for system management, in this example the random field estimation objective of Section 4.4.1 is used. This is defined as a special case of the VoI with the loss function of Equation 4-35, with **Z** defined in such a way as to apply a weight of one to estimates of the demand field and a weight of zero to estimates of the capacity field. Since measurements are available only of the demand, and the demand and capacity are independent, this definition only accounts for the prediction error in the demand field. The resulting optimized placements using this objective (along with the offline greedy algorithm) are indicated in Figure 6-15.

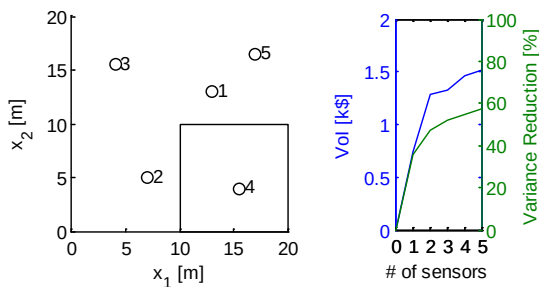

Figure 6-15: Placements for five humidity sensors optimized offline by the random field prediction metric (left) and the VoI and posterior random field variance reduction provided by these sensors (right).

As opposed to the sensors placed to optimize management in Figure 6-12, only one sensor is placed in the lower-right quadrant, as this objective does not include the economic incentive for sensing in this quadrant. Instead, placements are distributed more evenly across the system. Two sensors are placed in the upper-right quadrant, where the prior uncertainty in the demand field is highest. For comparison purposes, the VoI for these sensor placements is computed under the assumption of a cumulative system as in Section 6.2.1 in order to evaluate how sensors placed to improve random field estimation also help improve system management. In this example, these sensors achieved a VoI of $1520, 88% of that obtained by the placements of Figure 6-12. Placing sensors according to the prediction error metric in this case serves as a reasonable approach to approximately optimize sensor placements for management decision-making support under the less computationally efficient loss function for cumulative system management.

6.2.5. Comparison to Online Sensor Placement

Finally, online and offline sensor placement approaches to this problem are compared. In this example, it is assumed that observations of the local average humidity are obtained from each sensor shortly after it is placed, allowing for online optimization to be performed. The VoI formulation for a cumulative system as in Section 6.2.1 is again used, but VoI is optimized

following the online greedy approach described in Section 2.2.1.2. To demonstrate this online placement, the specific random instantiation of the system shown in Figure 6-10 is used. It should be kept in mind that in the online case the sequence of sensor placements depends on observations obtained from each previously placed sensor, and so the results of Figure 6-16 apply only to this instantiation of the random field.

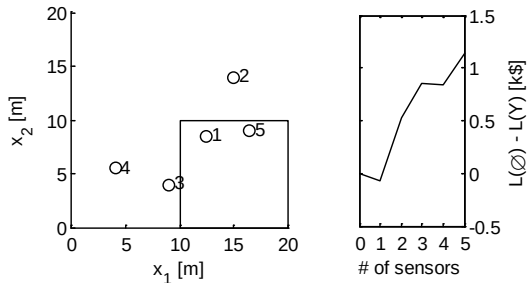

Figure 6-16: Placements for five humidity sensors optimized online by the VoI metric for a cumulative system (left) and the reduction in loss due to improved posterior actions (right).

While the first sensor placement is the same as was selected offline in Figure 6-12, placements differ after this due to updating of the prior model. To the right of Figure 6-16, the reduction in loss for managing this specific instantiation of the system after each sensor is added to the system (with optimal actions taken after measurements are obtained from this and all previously placed sensors) is plotted. Note that, while VoI, which represents the expected reduction in management costs, is always positive, actual management cost may be higher than the expected cost in any specific realization of the system. However, as more sensors are placed, there is a net reduction in management cost, and all five sensors provide a reduction of $1140 with respect to the cost incurred if no observations had been taken. This is contrasted with a reduction of $540 that would have been obtained for this instantiation of the system if the offline optimal sensor placement scheme indicated in Figure 6-12 were used.

The adaptivity gap, or the relative benefit of managing the system using online versus offline sensing, can be empirically quantified by comparing the performance of both approaches across multiple random instantiations of the system. Using 20 different simulations, and performing online sensing optimization for each case, cost reductions average to 31% (although this ranges from 0% to 100% depending on the specific instantiation). This is greater than the average reduction of 26% using the optimal placements selected offline. This demonstrates that, in the expected sense, online placements outperform offline ones by a small margin in this problem. A general proof that sensor placements selected online perform at least as well in an expected sense as those selected offline is provided in Appendix A. However, the degree to which online sensing outperforms offline sensing, i.e., the adaptivity gap, will vary on a problem-by-problem basis. Under the framework of adaptive submodularity (a generalization of submodularity to sequential sensing problems), guarantees on the performance of a greedy online sensor

placement approach can be provided (Golovin and Krause, 2011; Chen and Krause, 2013). Note that for some metrics, such as the random field estimation metric considered above, there is no adaptivity gap, as the value of the sensor placement is not affected by the collected measurement, and so there is no additional benefit to online placement (Krause and Guestrin, 2007).

Chapter 7

Extension to Spatio-Temporal Systems

This chapter presents an extension of the basic sensor placement methodologies discussed previously to spatio-temporal systems, i.e., systems in which the random field values vary across time as well as space. In principle, these systems can be handled as special cases of the general methodology presented previously, i.e., through the expansion of domain X to include a temporal dimension, and the appropriate definitions of random field variables in this dimension. However, in practice, several issues prevent this formulation from being easily applied. First, by increasing the dimensionality of the system into time, the exponential growth problem for the variable state space as discussed at the beginning of Chapter 5 is exacerbated. Furthermore, if actions for managing the system can be changed over time, this increases the dimensionality of the action choice space. Finally, if measurements within the system are associated with specific times, they will vary as the system evolves, and so the value of this information, and therefore the optimal sensor placements for managing the system, will be time-dependent. In addition, information propagation is unidirectional in time, i.e., only measurements taken prior to a given action being implemented can be used to help guide the choice of that action. These difficulties make sequential sensor placement and decision-making a computationally daunting problem.

The problem of inspection scheduling and sequential decision-making for system management has been investigated previously. For example, Straub has worked to develop efficient methods for risk-based inspection planning using optimal component-level strategies, along with information-sharing heuristics to extend these plans to system-level (e.g., Straub, 2004; Straub and Faber, 2005). Efficient methods for model updating for monitored deteriorating systems with multiple components based on hierarchical dynamic Bayesian networks are presented by Luque and Straub (2016). In addition, recent work by Memarzadeh has developed efficient methods to approximately allocate inspection efforts between components in a multi-component system, based on a hierarchical partially-observable Markov decision process model and a virtual-auction-based inspection allocation heuristic (e.g., Memarzadeh and Pozzi, 2015, 2016; Memarzadeh et al., 2016).

Instead of presenting an approach towards sensor placement and scheduling in spatio-temporal systems incorporating one or another of the above frameworks, this chapter instead focuses on extending the approaches developed for spatial systems in this work and identifying situations

where the VoI metric can be tractably computed, and thus where efficient sensor placement and scheduling can be addressed using the approaches developed in previous chapters.

It is assumed throughout this chapter that time is discretized and that the time horizon for system analysis is finite rather than infinite; these assumptions parallel those made for spatial systems, where the finite domain $\mathcal{X}$ of the system can be represented by a set of discrete coordinates X. Likewise, for this chapter, the temporal domain $\mathcal{T}$ is discretized to a finite set τ of n_τ time steps, such that the full spatio-temporal domain is represented by the set of discretized spatio-temporal coordinates $X \times \tau = \{(\mathbf{x}, t): \mathbf{x} \in X, t \in \tau\}$. Section 7.1 presents a formulation for VoI in spatio-temporal systems, efficient techniques for its computation, and a greedy approach to optimal sensor placement and scheduling, while Section 7.2 demonstrates the application of these methods in example problems.

Finally, note that the phrases *sensor placement*, *sensor scheduling*, and *sensor placement and scheduling* will have specific meanings in the context of this chapter. *Sensor placement* here refers to the selection of a set of measurements that are repeated for each time step. Intuitively, this corresponds to the selection of locations to place sensors, where these will collect data at regular intervals over their operational lifetime. *Sensor scheduling* here refers to the identification of times at which to collect a prescribed set of measurements on the system. This corresponds to a case where, once a number of sensors are deployed, they are interrogated at irregular intervals at the discretion of the managing agent. *Sensor placement and scheduling*, or simultaneous placement and scheduling, refers to the problem of jointly selecting which measurements to obtain and when to obtain them over the full spatio-temporal domain. This might correspond to a situation in which human or robotic inspectors are to be deployed within a system; the times at which they are to be deployed as well as the components they are to inspect at these times must both be determined.

7.1. Formulation of Spatio-Temporal Value of Information

This section presents the general formulation of the VoI metric and its use for optimal sensor placement, scheduling, and placement and scheduling in spatio-temporal systems. The basic modeling assumptions of Section 3.3 are relaxed in this section.

The random field affecting the spatio-temporal system is a function of space and time, defined as $\mathrm{f}(\mathbf{x}, t)$ over the space-time domain, where the temporal coordinate t is now explicitly separated from the spatial (or field indexing) coordinate $\mathbf{x}$. States, actions, and losses remain substantially the same as defined in Section 2.1, but with a somewhat different interpretation. For example, state vector $\mathbf{s}$ now describes the states of all system components at all discrete times in the time domain. While previously this vector may have consisted of n elements, one for each system component, in the corresponding spatio-temporal system it consists of nn_τ elements, describing the state of each component at each discrete time step. The limit-states and system states are likewise expanded to include different instantiations of the variables for each time step. The

system management action vector $\mathbf{a}$ and action set $\mathcal{A}$ may also be extended, e.g., if new actions are selected for each time step, or they may remain the same as in the purely spatial system, e.g., if the same actions are to be undertaken throughout the management time horizon. The definition of the loss function must be accordingly modified as well, to account for the total loss over the modeled management life-cycle of the system (perhaps discounted to a present value), taking into account all spatio-temporal random field variables or states and management actions over this lifetime.

As discussed in Sections 2.1.4 and 3.1, the assumption that management actions cannot affect the random field variables in the spatial system is preserved in the spatio-temporal system model. This is again done to ensure the acyclic nature of the PGM. However, in the case of a temporally changing system, this assumption is somewhat more difficult to enforce. In general, management activities in one time step will affect the behavior of the system in future time steps. In Markov process models and dynamic Bayesian networks, the creation of separate instantiations of random field variables for each time step ensures the acyclic nature of the model. An alternative approach, which is adopted here, it to capture these effects within the loss function. An arbitrarily defined loss function $\mathrm{L}(\mathbf{f}, \mathbf{a})$ of all random variables and all management actions can theoretically account for any action effects on the evolution of the system, either directly where these effects are deterministic, or indirectly through augmentation of vector $\mathbf{f}$ where these effects are stochastic. In practice, however, the appropriate definition of such a loss function will be difficult, and so the approach of using separate variable instantiations at each time step and employing a dynamic Bayesian network or Markov process analysis approach will be attractive for most practical applications, using methods such as those discussed at the beginning of this chapter. This section will instead focus on first presenting a general formulation for the VoI in systems where such an approach is not followed, and then examining the specific case of a temporally cumulative system, where this approach can be applied easily and efficiently using the methods of previous chapters.

In spatio-temporal systems, VoI is still formulated as in Equation 4-6 as the difference of the prior and posterior expected loss. Furthermore, the prior expected loss is still computed as in Equation 4-1; it remains the minimum over all choices of management actions over the time horizon for system management of the expected value of the loss function across all of the spatio-temporal random field variables. Since no new information will be available at any point, there is no need to consider decision-making for this system sequentially; an optimal plan can be determined at the beginning of the management process and carried out over time.

In the case of the posterior expected loss of Equation 4-5, the formulation must be modified somewhat, as the collection of information and the selection of optimal posterior actions based on this information is carried out in sequence. To capture these effects, let the set of measurements Y be expressible as $Y = \bigcup_{t=1}^{n_\tau} Y^{(t)}$, with the observation vector $\mathbf{y}$ likewise expressible as $\mathbf{y} = \{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(n_\tau)}\}$. This notation is used to denote the fact that information is collected over time, with measurement set $Y^{(t)}$ corresponding to measurements made at time

index t, and observations vector $\mathbf{y}^{(t)}$ corresponding to the outcome of these measurements. Furthermore, the vector of all measurement information obtained up to and including time index t is denoted $\mathbf{y}^{-(t)} = \{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(t)}\}$, concatenating measurement results of $Y^{-(t)} = Y^{(1)} \cup \dots \cup Y^{(t)}$. Note that for this purpose, the nominal information collection time index does not necessarily refer to when a measurement is taken, or the time index for the random variables that are being measured, but rather to the earliest time index that this information is available to the managers of the system to support decision-making. Also note that some of the sets and observations associated with certain times may be empty, if no information is to be collected at that time according to the inspection plan.

Similarly, action vector $\mathbf{a}$ is expressed as $\mathbf{a} = \{\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(n_\tau)}\}$, where $\mathbf{a}^{(t)}$ refers to a subset of actions that are chosen at time index t. As with observations, although these actions may take effect later than this, t is the time index at which these actions are selected and locked, i.e., they are fixed and cannot be changed later. The set of possible actions at time index t is denoted $\mathcal{A}^{(t)}$. Note also that the action vectors and potential action sets associated with a time index may be empty if no action choices are associated with that time step.

Referencing Equation 4-4, the posterior expected loss of Equation 4-5 can be expressed equivalently as:

$$\mathbb{E}\mathrm{L}(Y) = \mathbb{E}_Y \mathbb{E}_{F|\mathbf{y}} \mathrm{L}\big(\mathbf{f}, \mathbf{a}^*(\mathbf{y})\big) \tag{7-1}$$

This formulation also applies in the spatio-temporal case, where the optimal action set $\mathbf{a}^*$ taking measurements $\mathbf{y}$ into account denotes the actual sequence of optimal actions chosen to manage the system over the given time horizon. However, note that the optimal actions are not all selected simultaneously; instead, decision-making is made sequentially for each time step, taking into account only information gathered up to that time step. Thus, the set of optimal actions is expressed as:

$$\mathbf{a}^*(\mathbf{y}) = \{\mathbf{a}^{*(1)}\big(\mathbf{y}^{-(1)}\big), \mathbf{a}^{*(2)}\big(\mathbf{y}^{-(2)}\big), \dots, \mathbf{a}^{*(n_\tau)}\big(\mathbf{y}^{-(n_\tau)}\big)\} \tag{7-2}$$

This set of optimal actions is built sequentially over the management time horizon of the system. At any time step t, an optimal plan is made for the management of the system at that and all future time steps, given all information collected up to that time step and considering the optimal actions taken in all previous time steps:

$$\mathbf{a}^{*(t)}\big(\mathbf{y}^{-(t)}\big) = \\ \operatorname{argmin}_{\mathbf{a}^{(t)} \in \mathcal{A}^{(t)}} \mathbb{E}_{Y^{(t+1)}|\mathbf{y}^{-(t)}} \min_{\mathbf{a}^{(t+1)} \in \mathcal{A}^{(t+1)}} \cdots \mathbb{E}_{Y^{(n_\tau)}|\mathbf{y}^{-(n_\tau-1)}} \min_{\mathbf{a}^{(n_\tau)} \in \mathcal{A}^{(n_\tau)}} \mathbb{E}_{F|\mathbf{y}^{-(n_\tau)}} \mathrm{L}(\mathbf{f}, \mathbf{a}) \tag{7-3}$$

where $\mathbf{a}^{*(t)}$ denotes the optimal actions for time step t. The sequence of all actions in this equation is $\mathbf{a} = \{\mathbf{a}^{*-(t-1)}, \mathbf{a}^{(t)}, \mathbf{a}^{+(t+1)}\}$, where $\mathbf{a}^{+(t+1)}$ denotes actions chosen to manage the system from time step $t+1$ forward, chosen from potential action set $\mathcal{A}^{+(t+1)}$, and $\mathbf{a}^{*-(t-1)} =$

$\{\mathbf{a}^{*(1)}(\mathbf{y}^{-(1)}), \dots, \mathbf{a}^{*(t-1)}(\mathbf{y}^{-(t-1)})\}$ denotes the set of optimal actions selected previous to time step t. The optimized actions corresponding with the current time step $\mathbf{a}^{*(t)}(\mathbf{y}^{-(t)})$ is then implemented, becoming part of the optimal action set $\mathbf{a}^{*-(t)}$, and the process is repeated for the next time step, using the next set of available data $\mathbf{y}^{-(t+1)}$. By repeating this process of information collection, optimization, and implementation of the next group of optimal actions at each time step, the set of implemented optimal management actions $\mathbf{a}^*$ conditional to observation $\mathbf{y}$ is built up. The result allows the conditional expected loss $\mathbb{E}\mathrm{L}(\mathbf{y})$ to be computed. This process must then be repeated for different realizations of $\mathbf{y}$ such that the expectation over Y can be computed, as shown in Equation 7-1.

This sequential decision-making adds an additional level of complexity over that of the computation of VoI in static systems as discussed in Chapter 5. This compounds the complexity introduced by expanding the number of possible variable states, observations, and actions in accordance with the expansion of the system in the temporal dimension (e.g., a transition from n state variables to nn_τ state variables). For this reason, in a general case, analysis of VoI in this manner is infeasible. More efficient approaches, such as those discussed at the beginning of this chapter, should be attempted. However, motivated by the spatial decomposability of the loss function of Equation 3-1 and the resulting improvement in computational efficiency discussed in Section 5.2, the next subsection presents the special case of a temporally decomposable loss function, under which spatio-temporal VoI can be efficiently computed.

7.1.1. Temporally Decomposable Loss Functions

Assume that the loss function for a spatio-temporal system decomposes as follows:

$$\mathrm{L}(\mathbf{s}, \mathbf{a}) = \sum_{t=1}^{n_\tau} \gamma(t) \mathrm{L}^{(t)}(\mathbf{s}^{(t)}, \mathbf{a}^{(t)}) \tag{7-4}$$

where $\gamma(t)$ is a discounting factor. Under this assumption, the total loss is equal to the discounted sum of losses for managing the system at each time step, where this loss is a function only of the state of the system and the actions taken for its management at this time. Note that, as in Equation 3-1, the loss function is expressed as a function of state vector $\mathbf{s}^{(t)}$ for time index t; this loss might also be expressed as a function of $\mathbf{f}^{(t)}$, $\mathbf{g}^{(t)}$, or $\mathbb{s}^{(t)}$.

Following the derivation of Equation 5-4, the posterior expected loss in this case can be expressed as:

$$\mathbb{E}\mathrm{L}(Y) = \sum_{t=1}^{n_\tau} \gamma(t) \mathbb{E}_{Y^{-(t)}} \min_{\mathbf{a}^{(t)} \in_c \mathcal{A}^{(t)}} \mathbb{E}_{S^{(t)}|\mathbf{y}^{-(t)}} \mathrm{L}^{(t)}(\mathbf{s}^{(t)}, \mathbf{a}^{(t)}) = \sum_{i=1}^{n} \gamma(t) \mathbb{E}\mathrm{L}^{(t)}(Y^{-(t)}) \tag{7-5}$$

Note here that the outer expectation is only performed with respect to $Y^{-(t)}$, as only past measurements can be used to optimize the choice of action at time step t. Also note that the minimization is performed with respect to actions $\mathbf{a}^{(t)}$, as past and future actions have no effect on the loss function for time step t. Finally, note that the inner expectation is performed with

respect to states $\mathbf{s}^{(t)}$ at time index t, conditional to previous measures $\mathbf{y}^{-(t)}$, as only these states affect the temporally local loss function. VoI can therefore be decomposed, similar to Equation 5-5, as:

$$\mathrm{VoI}(Y) = \sum_{t=1}^{n_\tau} \gamma(t)\, \mathrm{VoI}^{(t)}\left(Y^{-(t)}\right) \tag{7-6}$$

This temporal decomposition of the VoI is analogous to the spatial decomposition discussed in Section 5.2, and provides comparable computational advantages. Note, however, that for any time step, the temporally local expected loss may itself be difficult to compute, e.g., due to a large number of possible joint states $\mathbf{s}^{(t)}$, actions $\mathbf{a}^{(t)}$, and prior observations $\mathbf{y}^{-(t)}$.

If, in addition to being temporally decomposable, the loss function is also spatially decomposable, then the loss is expressible as a nested sum over time steps and over the system components as:

$$\mathrm{L}(\mathbf{s}, \mathbf{a}) = \sum_{t=1}^{n_\tau} \gamma(t) \sum_{i=1}^{n} \mathrm{L}_i^{(t)}\left(\mathbf{s}_i^{(t)}, \mathbf{a}_i^{(t)}\right) \tag{7-7}$$

Such a spatio-temporally decomposable loss function allows the temporal decomposition of the VoI discussed in this section, as well as its spatial decomposition as in Section 5.2, combining the advantages of both approaches.

It is important to note that temporal decomposability of the loss function in Equation 7-4 implies that actions taken at a certain time step do not affect the system in future time steps, e.g., that past actions have no effect on the future behavior of the system. While this assumption has beneficial computational implications, it will not be generally valid. Appropriate definition of the system management problem in an approximate manner, such that the resulting loss function is temporally decomposable, may therefore be attractive for reasons of computational efficiency.

7.1.1.1. Relationship to the Bellman Equation

For many sequential decision-making problems, especially in Markov decision processes, evaluation of VoI is performed by means of the Bellman Equation, a dynamic programming approach (Bellman, 1984). This subsection examines the relationship between this approach and the discussions of this section relating to the spatio-temporal VoI. For the purposes of this subsection, it is assumed that the loss function is temporally decomposable as in Equation 7-4, but that the evolution of random field variables F is affected by the choice of management actions, i.e., these coupling effects are no longer captured as part of the loss function formulation. Choices of actions will now affect the future evolution of the system, coupling the optimization of actions between time steps. This is a common assumption in Markov decision processes, where, for first-order processes, the state of the system at a time step is a random function of its state at the previous time step and the actions taken to manage the system at the previous time step. In this subsection, however, no such Markovian assumption is made.

Within this sequential decision-making problem, a dynamic programming approach to VoI evaluation is adopted. This approach makes use of a value function $\mathrm{V}^{(t)}\left(\mathbf{a}^{-(t)}, \mathbf{y}^{-(t)}\right)$ of all actions and measurements taken up to and including that time (Bellman, 1984). The value function is also sometimes known as the "cost to go" function, as it represents the minimum expected loss to manage the system from time t forward. This value function is defined recursively for a temporally decomposable loss function as (Memarzadeh and Pozzi, 2016):

$$\mathrm{V}^{(t)}\left(\mathbf{a}^{-(t)}, \mathbf{y}^{-(t)}\right) = \mathbb{E}_{S^{(t)}|\mathbf{a}^{-(t)},\mathbf{y}^{-(t)}} \mathrm{L}^{(t)}\left(\mathbf{s}^{(t)}, \mathbf{a}^{(t)}\right) + \frac{\gamma(t+1)}{\gamma(t)} \mathbb{E}_{Y^{(t+1)}|\mathbf{a}^{-(t)},\mathbf{y}^{-(t)}}\left[\min_{\mathbf{a}^{(t+1)}\in\mathcal{A}^{(t+1)}} \mathrm{V}^{(t+1)}\left(\mathbf{a}^{-(t+1)}, \mathbf{y}^{-(t+1)}\right)\right] \quad (7\text{-}8)$$

That is, the value function is the expected loss at the current time step plus the discounted expected value (over possible observations at the next time step) of the minimum (across management actions taken at the next time step) of the value function at the next time step. Following this recursive definition, the posterior expected loss can be evaluated as:

$$\mathbb{E}\mathrm{L}(Y) = \mathbb{E}_{Y^{(1)}} \min_{\mathbf{a}^{(1)}\in\mathcal{A}^{(1)}} \mathrm{V}^{(1)}\left(\mathbf{a}^{-(1)}, \mathbf{y}^{-(1)}\right) \quad (7\text{-}9)$$

noting that $\mathbf{a}^{-(1)} = \mathbf{a}^{(1)}$ and that $\mathbf{y}^{-(1)} = \mathbf{y}^{(1)}$. Prior expected losses can be evaluated similarly by dropping any conditioning on observations from Equations 7-8 and 7-9.

Computation of the VoI in this manner is in general computationally demanding for the same reasons discussed previously, namely the need to account for all potential sequences of prior observations and actions when evaluating the value function. However, under the assumption that actions do not affect the evolution of the underlying random field affecting the system, the value function can be expressed as:

$$\mathrm{V}^{(t)}\left(\mathbf{y}^{-(t)}\right) = \min_{\mathbf{a}^{(t)}\in\mathcal{A}^{(t)}} \left[\mathbb{E}_{S^{(t)}|\mathbf{y}^{-(t)}} \mathrm{L}^{(t)}\left(\mathbf{s}^{(t)}, \mathbf{a}^{(t)}\right)\right] + \frac{\gamma(t+1)}{\gamma(t)} \mathbb{E}_{Y^{(t+1)}|\mathbf{y}^{-(t)}} \mathrm{V}^{(t+1)}\left(\mathbf{y}^{-(t+1)}\right) \quad (7\text{-}10)$$

where all conditioning of states and observations on past actions has been removed and the minimization step has been moved through the iteration, e.g., minimization over $\mathbf{a}^{(t)}$ is performed as the outermost operation as part of iteration t rather than the minimization over $\mathbf{a}^{(t+1)}$ performed as an inner operation in Equation 7-8. Note that the value function is now independent of the sequence of prior actions and that, as in Equation 7-5, the value function can be evaluated as the discounted sum of losses for each time step due to the distributive property of the expectation. This shows that the formulation of VoI with temporally decomposable loss functions discussed above is compatible with a sequential decision process within the Bellman Equation framework under the assumption that actions do no effect the evolution of the system state.

7.1.2. Offline and Online Sensor Placement and Scheduling

Sensor placement and/or scheduling in spatio-temporal systems can be understood in the context of the general optimization objective of Equation 2-16, where the set Y describes the complete sensing plan in space and time (e.g., where to place sensors, and when to collect measurements from these sensors). The problems of sensor placement, scheduling, and simultaneous placement and scheduling can then be understood as representing different constraint functions $\mathrm{C}(Y)$ on the sensing objective. For example, in a sensor placement problem, there would be no additional cost associated with selecting a new measurement at a previously selected location but at a different time, while there is an additional cost to choosing a new measurement at a previously unmeasured location. An optimal solution would therefore include measurements across all time steps for a given sensor, because the benefits of doing so, even if they were small, would always outweigh the cost. Furthermore, for sensor scheduling, including a measure at the same time as other measures will incur no (or negligible) additional cost, while adding a new measurement at a time when no other measurements are specified increases cost greatly. Finally, for simultaneous placement and scheduling, the cost might for example be linear with respect to the number of measurements, regardless of their locations in space or time. For measurements not associated with a particular location in space, or measurements that take place across multiple time steps, the cost function would have to be defined accordingly. Optimized placement and scheduling is a typical problem in the deployment of wireless sensing networks, where connectivity issues and power requirements must be considered (e.g., Meliou et al., 2007; Krause et al., 2011).

For optimizing sensing in spatio-temporal systems, offline optimization is occasionally considered, for example when sensors are to be installed in a new structure and their configuration cannot be easily changed. Approximate offline optimization using the greedy algorithm approach can be implemented in spatio-temporal systems following the approach of Section 2.2.1.1, with appropriate use of the spatio-temporal VoI metric. However, in general, online optimization is a natural approach to take when measuring spatio-temporal systems. Furthermore, the unidirectional nature of time imposes a natural sequence on the observations of the system. The online optimization approach of Section 2.2.1.2 must therefore be slightly modified in the case of spatio-temporal systems.

The modified approach can be described as a nested optimization problem. At each time step, using all information gathered up to that time step, an offline sensor placement approach is adopted to optimize the set of measurements obtained for current and future time steps. Then, the subset of measurements within this optimal set corresponding to the current time step is implemented, providing new measurements to help guide future optimal sensor placements in the next time step and onwards. This approach is outlined in Algorithm 2, where the use of the VoI metric for sensing is assumed. Note that the inner sensing optimization problem described in this algorithm can be solved using the offline optimization approach of Algorithm 1 and Section 2.2.1.1.

Algorithm 2: Pseudo-code for online sensing optimization in spatio-temporal systems.

Input candidate measurement set $\mathcal{Y}$, candidate action set $\mathcal{A}$, objective function $\mathrm{M}_{\mathrm{VoI}}(\cdot)$, cost function $\mathrm{C}(\cdot)$, constraint b

$Y^{*-(0)} = \emptyset$, $\mathbf{y}^{-(0)} = \emptyset$, $\mathcal{Y}^{+(1)} = \mathcal{Y}$, $\mathcal{A}^{+(1)} = \mathcal{A}$, $\mathbf{a}^{*-(0)} = \emptyset$

for $t = 1, \dots, n_\tau$

$Y^{(t)} \cup \dots \cup Y^{(n_\tau)} = \mathrm{argmax}_{Y^{+(t)} \subseteq \mathcal{Y}^{+(t)}} \ \mathrm{M}_{\mathrm{VoI}}\left(Y^{+(t)} | \mathbf{y}^{-(t-1)}, \mathbf{a}^{*-(t-1)}\right)$

subject to $\mathrm{C}\left(Y^{*-(t-1)} \cup Y^{(t)} \cup \dots \cup Y^{(n_\tau)}\right) \leq b$

measure $Y^{(t)}$ **to obtain** $\mathbf{y}^{(t)}$

$\mathbf{y}^{-(t)} = \left\{\mathbf{y}^{-(t-1)}, \mathbf{y}^{(t)}\right\}$

$$\mathbf{a}^{*(t)}\left(\mathbf{y}^{-(t)}\right) = \mathrm{argmin}_{\mathbf{a}^{(t)} \in \mathcal{A}^{(t)}} \ \mathbb{E}_{Y^{(t+1)}|\mathbf{y}^{-(t)}} \min_{\mathbf{a}^{(t+1)} \in \mathcal{A}^{(t+1)}} \cdots \mathbb{E}_{F|\mathbf{y}^{-(n_\tau)}} \mathrm{L}(\mathbf{f}, \mathbf{a}) \quad \text{(Eqn. 7-3)}$$

$\mathbf{a}^{*-(t)} = \left\{\mathbf{a}^{*-(t-1)}, \mathbf{a}^{*(t)}\left(\mathbf{y}^{-(t)}\right)\right\}$

$Y^{*-(t)} = Y^{*-(t-1)} \cup Y^{(t)}$

$\mathcal{Y}^{+(t+1)} = \mathcal{Y}^{+(t)} \backslash \mathcal{Y}^{(t)}$

$\mathcal{A}^{+(t+1)} = \mathcal{A}^{+(t)} \backslash \mathcal{A}^{(t)}$

end

Output $Y^{*-(n_\tau)}$

The VoI metric for online optimal sensing used here, similar in form to Equation 4-13, is:

$$\mathrm{M}_{\mathrm{VoI}}\left(Y^{+(t)} | \mathbf{y}^{-(t-1)}, \mathbf{a}^{*-(t-1)}\right) = \mathrm{VoI}\left(Y^{+(t)} | \mathbf{y}^{-(t-1)}, \mathbf{a}^{*-(t-1)}\right) - \mathrm{C}\left(Y^{*-(t-1)} \cup Y^{+(t)}\right) + \mathrm{C}\left(Y^{*-(t-1)}\right) \quad (7\text{-}11)$$

This represents the difference between the VoI of future measurements conditional to the outcomes of past measurements and past actions and the additional cost of including future measurements. The conditional VoI is:

$$\mathrm{VoI}\left(Y^{+(t)} | \mathbf{y}^{-(t-1)}, \mathbf{a}^{*-(t-1)}\right) = \mathbb{E}\mathrm{L}\left(\mathbf{y}^{-(t-1)}, \mathbf{a}^{*-(t-1)}\right) - \mathbb{E}\mathrm{L}\left(Y^{+(t)} | \mathbf{y}^{-(t-1)}, \mathbf{a}^{*-(t-1)}\right) \quad (7\text{-}12)$$

where:

$$\mathbb{E}\mathrm{L}\left(\mathbf{y}^{-(t-1)}, \mathbf{a}^{*-(t-1)}\right) = \min_{\mathbf{a}^{+(t)} \in \mathcal{A}^{+(t)}} \mathbb{E}_{F|\mathbf{y}^{-(t-1)}} \mathrm{L}\left(\mathbf{f}, \left\{\mathbf{a}^{*-(t-1)}, \mathbf{a}^{+(t)}\right\}\right) \quad (7\text{-}13)$$

and:

$$\mathbb{E}\mathrm{L}\left(Y^{+(t)} | \mathbf{y}^{-(t-1)}, \mathbf{a}^{*-(t-1)}\right) = \mathbb{E}_{Y^{+(t)}|\mathbf{y}^{-(t-1)}} \mathbb{E}_{F|\mathbf{y}} \mathrm{L}\left(\mathbf{f}, \left\{\mathbf{a}^{*-(t-1)}, \mathbf{a}^{*+(t)}(\mathbf{y})\right\}\right) \quad (7\text{-}14)$$

where $\mathbf{a}^{*+(t)}(\mathbf{y})$ is computed following the sequential approach of Equations 7-2 and 7-3. Note that this greedy approach remains myopic, as the benefit of information with respect to the possibility of improving future sensor placements is not taken into account. Also note that this approach has been formulated generally, i.e., not assuming temporal decomposability of the loss

function. Temporal decomposability allows for easier evaluation of the expected losses and optimization of actions for each time step.

7.2. Demonstrative Examples

This section details the application of the above methodology for optimal sensing in spatio-temporal systems through three examples. Section 7.2.1 presents an example to binary classification in one spatial dimension to demonstrate a trivial application of the method, with intuitive outcomes. Section 7.2.2 presents an example in one spatial dimension using a series system topology, demonstrating a case where the loss function decomposes only temporally. Section 7.2.3 presents an example of a structural management problem to illustrate how the method can be applied for infrastructures.

7.2.1. Classification in an Evolving 1-D Random Field

This section demonstrates optimal offline spatio-temporal sensor placement and scheduling for binary classification of a temporally varying random field in one spatial dimension. A regular grid of spatial coordinates, with spacing of 0.2 distance units between grid points, defines the discretized spatial domain X of the problem, with $n_X = 25$. Five time steps, spaced at unit time, define the discretized time horizon τ, with $n_\tau = 5$.

A random field is defined over this spatio-temporal domain by a Gaussian process model in time and space. Following the notation of a spatial Gaussian process in Equation 3-2, the spatio-temporal process is denoted:

$$\mathrm{f}(\mathbf{x}, t) \sim \mathcal{GP}\left(\mu(\mathbf{x}, t), \mathrm{k}\left(\mathbf{x}_i, t_i, \mathbf{x}_j, t_j\right)\right) \tag{7-15}$$

where $\mu(\mathbf{x}, t)$ is the spatio-temporal mean function and $\mathrm{k}\left(\mathbf{x}_i, t_i, \mathbf{x}_j, t_j\right)$ is the spatio-temporal covariance. A typical assumption, which is made for this section and elsewhere in this work, is that the spatio-temporal covariance is separable as the product of spatial and temporal covariance functions (Cressie and Wikle, 2011):

$$\mathrm{k}\left(\mathbf{x}_i, t_i, \mathbf{x}_j, t_j\right) = \mathrm{k}_X\left(\mathbf{x}_i, \mathbf{x}_j\right) \mathrm{k}_\tau\left(t_i, t_j\right) \tag{7-16}$$

In this example, a zero-mean field is assumed, i.e., $\mu(\mathbf{x}, t) = 0$. The spatial and temporal covariance are separately defined with unit standard deviation throughout the spatio-temporal domain and using the square exponential correlation function of Equation 6-3, with the spatial correlation length scale being two spatial units and the temporal correlation time scale being two temporal units. Measurements are possible at each of the 125 discrete spatio-temporal coordinates, with uncorrelated noise between measurements and standard deviation of the noise of 0.1. Local limit-states are defined following Equation 3-11 with $\mathbf{\Xi}_G = \mathbf{I}$ and $\mathbf{g}_0 = \mathbf{0}$, and $n = 125$ binary states are defined following Equation 3-14 for each space-time coordinate. In

this example, the misclassification loss function of Equation 4-37 is used, resulting in a spatio-temporally decomposable loss function as in Equation 7-7. No discounting is considered, i.e., $\gamma(t) = 1$.

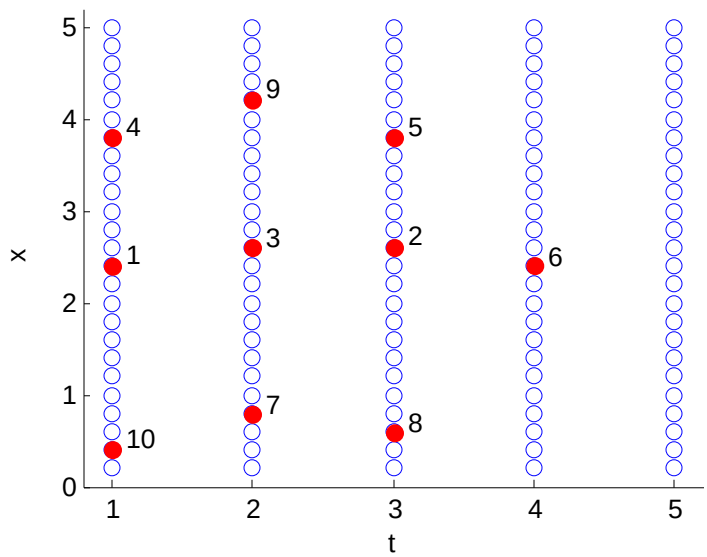

Figure 7-1: Optimal placement and scheduling of ten measurements in a one-dimensional system for classification.

Placements and scheduling of ten measurements following offline greedy optimization are indicated in Figure 7-1, with empty circles indicating potential measures and filled circles indicating selected measures, along with the order of their selection. Note the systematic, symmetric approach to sensing. As there is nothing to distinguish one measure from another initially, except for the time difference, the first measurement is prescribed for the first time slice in the center of the domain, such that this measure can inform decision-making over the spatial domain and at later time steps. The second measure is prescribed for the center of the space and time domains. Subsequent measures are spaced around these points, both in space and time. Prior expected loss for this problem is 62.5, i.e., half of the spatio-temporal locations are expected to be misclassified a priori. With these ten measurements (less than a tenth of all possible measurements), posterior expected loss is about 45, indicating that nearly two-thirds of the states will be correctly predicted a-posteriori.

7.2.2. Management of a River Water Intake System

This section presents an example motivated by a water intake system along a river, and demonstrates the application of the approaches developed in this chapter to a system whose loss function is decomposable in time but not in space. Consider a 10km stretch of a river, along

which the water contamination level (measured in ppm of an unspecified contaminant) varies in space and time, as water flows down the river. Along this spatial domain, $n = 5$ water intake stations divert river water for municipal use. If the contamination level at any of these stations exceeds a threshold of 100ppm, the water supply system will be contaminated, and a cost of $C_f = \$10\text{k}$ is incurred for every hour that contaminated water is drawn into an intake. To prevent this, a reserve water supply can be substituted for any one of the intakes, at a cost of $C_r = \$1\text{k}$ per hour per intake.

The contamination level at the five intake locations over a five-hour period (discretized to one-hour intervals) is modeled as a log-Gaussian random field, such that under the logarithm transformation, the field is characterized by a Gaussian process model, as discussed in Section 3.3. The mean function of the Gaussian process is defined as $\mu(\mathbf{x}, t) = 2.3$, corresponding to a median contamination level of 10ppm. The covariance function of the Gaussian process is:

$$\mathrm{k}(x_i, t_i, x_j, t_j) = \sigma^2 \exp\left(-\frac{\left((x_i - vt_i) - (x_j - vt_j)\right)^2}{2\lambda^2}\right) \tag{7-17}$$

with $\sigma = 1.3$ (corresponding to 50 ppm standard deviation), $v = 1$ kph, and $\lambda = 1$ km. This function uses a modified version of the square exponential correlation function of Equation 6-3, including the effect of water flowing downriver at a constant velocity v. Thus, the contaminant concentration at a certain location will be strongly correlated with those of nearby locations, but also with those of upstream locations at earlier times and of downstream locations at later times. Diffusion of the contaminants is ignored in this problem; concentrations only vary as the water flows past the intake sites.

Limit-states and local states are defined as in Equations 3-11 and 3-14, with $\mathbf{\Xi}_G = -\mathbf{I}$ and $\mathbf{g}_0 = \mathbf{1}\log(100\text{ppm})$. Measurements are considered to be possible at each intake station at each hour. For consistency with the log-Gaussian contamination levels, these measurements are considered to have multiplicative log-Gaussian noise with median 1 and standard deviation 0.01, such that following the logarithm transformation the noise will be zero-mean additive Gaussian noise, as in Equation 3-5. For the management of this system, at each time step a prototypical loss function for a series system is considered, following Equation 4-17, with costs as outlined above. Due to the short time horizon, no discounting is considered, i.e., $\gamma(t) = 1$.

Figure 7-2 shows the results for optimal offline placement and scheduling of five contamination measures, with empty circles indicating potential measures (with the spatial coordinates of these being the locations of the intakes with respect to their distance downstream from the first intake), filled circles indicating selected measures with the order in which they are greedily selected indicated, and the background shading showing how these measures reduce the posterior uncertainty in the contamination field. Note that the measurements prescribed at earlier times at upstream locations will reduce the uncertainty in contamination levels not only at the location and time they are made, but also for downstream locations at later times. The indicated

measurement plan provides a VoI of \$4k, compared to a VoI for the first greedily selected measure of \$0.9k, and a VoPI for measuring contamination at all intakes at all times of \$7.5k. Therefore, the proposed scheme, which provides more than half of the maximum VoI while measuring only one fifth of possible locations, provides a reasonable quality solution to this spatio-temporal sensing problem for system management.

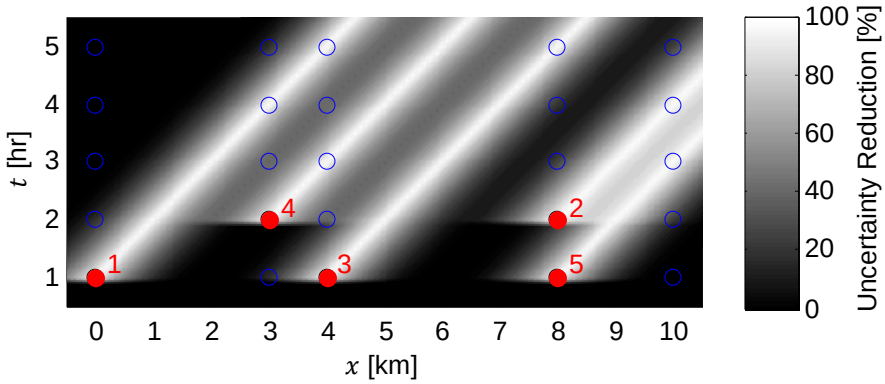

Figure 7-2: Optimal placement and scheduling of five water contamination samples along a 10km stretch of river to support water intake system management, plotted above the contamination level uncertainty reduction.

7.2.3. Differential Settlement in a Foundation System

This example application is motivated by the problem of differential settlement between columns of a structure (e.g., Glisic et al., 2005). The total settlements under each of nine columns of a structure over a ten-year period are modeled using a Gaussian process. The mean for this process is decreasing with time, modeling the average settlement of the columns:

$$\mu(\mathbf{x}, t) = 0.4\text{m}\left(1 - \exp\left(-\frac{t-1}{5\text{yr}}\right)\right) \tag{7-18}$$

The covariance function for the settlement is separable, as in Equation 7-16, with:

$$\mathrm{k}_X(\mathbf{x}_i, \mathbf{x}_j) = 0.2\exp\left(-\frac{(\mathbf{x}_i-\mathbf{x}_j)^{\mathrm{T}}(\mathbf{x}_i-\mathbf{x}_j)}{2(50\text{m})^2}\right) + 0.8\frac{\mathbf{x}_i{}^{\mathrm{T}}\mathbf{x}_j}{(50\text{m})^2} \tag{7-19}$$

and:

$$\mathrm{k}_\tau(t_i, t_j) = \left(\frac{0.005\text{m}}{\text{yr}}\right)^2 t_i\, t_{\mathrm{j}} \exp\left(-\frac{|t_i-t_j|^2}{2(5\text{yr})^2}\right) \tag{7-20}$$

This covariance structure describes, in space, the weighted average of a square exponential correlation function and a linear correlation function, with the linear portion tending to constrain

the settlement values to fall on a plane, while the square exponential portion introduces a perturbation to these co-planar settlements. In time, variance of the settlements increases with time, i.e., the total settlement is less certain farther into the future, and the correlations of settlements at different times is square exponential.

Measurements of each column's settlement are possible at the end of each year, with independent zero-mean noise with a standard deviation of 0.01m. These measurements are available to support decision-making beginning with the year after they are collected. Differential settlement of the columns is defined with respect to the difference between the settlement of each individual column and the average settlement of all columns, with the limit-states defined as in Equation 3-11 with $\boldsymbol{\Xi}_G$ as in Equation 5-32 and $\mathbf{g}_0 = \mathbf{0}$. Note that this leads to only n limit-state varialbes for the system, as opposed to $n(n-1)/2$ if the pairwise differential settlements between columns were considered. Component states are defined relating to whether or not the absolute differential settlement for each column exceeds a threshold of 0.05m, i.e., for each of $n = 9$ columns, the state at time t is defined as:

$$s_i^{(t)} = \mathbb{I}\left(g_{\text{upper}} \geq g_i^{(t)} \geq g_{\text{lower}}\right) \tag{7-21}$$

where $g_{\text{upper}} = 0.05\text{m}$ and $g_{\text{lower}} = -0.05\text{m}$. Note that, although this differs from the binary state definition of Equation 3-14, this still defines a binary state for each component, and the efficient techniques of Section 5.2.2 can still be applied. In particular, in Equation 5-24, instead of substituting (in the case of a spatio-temporal system) $P_{f,i}^{(t)} = \Phi\left(-\beta_i^{(t)}\right)$, the following substitution is instead made:

$$P_{f,i}^{(t)} = 1 - \Phi\left(\frac{g_{\text{upper}}}{\sigma_{g_i^{(t)}}} - \beta_i^{(t)}\right) + \Phi\left(\frac{g_{\text{lower}}}{\sigma_{g_i^{(t)}}} - \beta_i^{(t)}\right) \tag{7-22}$$

Following this definition, the probability of failure will correspond to the state definition of Equation 7-21. Note that this expression is still only a function of $\beta_i^{(t)}$, and so Equation 5-24 can still be evaluated using only a univariate expectation over $\beta_i^{(t)}$, where the distribution of $\beta_i^{(t)}$ is defined for component i at time t following the procedure outlined in Section 5.2.2.

Management of the system is defined using a spatio-temporally decomposable loss function as in Equation 7-7, with the loss for each time step described using the prototypical cumulative system loss function of Equation 4-15, with $C_{f,i} = \$100\text{k}$ and $C_{r,i} = \$20\text{k}$ for all columns. That is, if no action is taken and the difference of a column's settlement from the average exceeds the threshold, a $\$100\text{k}$ penalty is incurred for that column in that year, in terms of damage to the structure due to the uneven settlement that must be remedied. If excessive settlements are expected, the managing agent has the option of taking preventative actions at a cost of $\$20\text{k}$ per column per year, which will prevent damages due to differential settlement. A discounting factor of 10% per year is considered, i.e.:

$$\gamma(t) = 0.9^{(t-1)} \tag{7-23}$$

Figure 7-3a plots the spatial locations of the nine columns. Figure 7-3b shows the results of offline optimal sensor placement and scheduling for measuring settlement of these columns. This scheme focuses on the corner columns (1, 3, 7, and 9), measuring these continuously after the third year. Temporally, later measures are prioritized, as failures are more likely then due to the higher variance. However, no measurements are made in the tenth year, as the outcomes of these measures will be available too late to support decision-making. This scheme, with 37 measures, maximizes the net VoI, as shown in Figure 7-3c which illustrates the trends in VoI and net VoI (assuming a cost of $1k per measure, discounted to present value using the same discount rate as the loss) provided by greedily selected schemes with increasing numbers of measurements. Note that the increase in VoI obtained with each additional sensor is diminishing as more sensors are added.

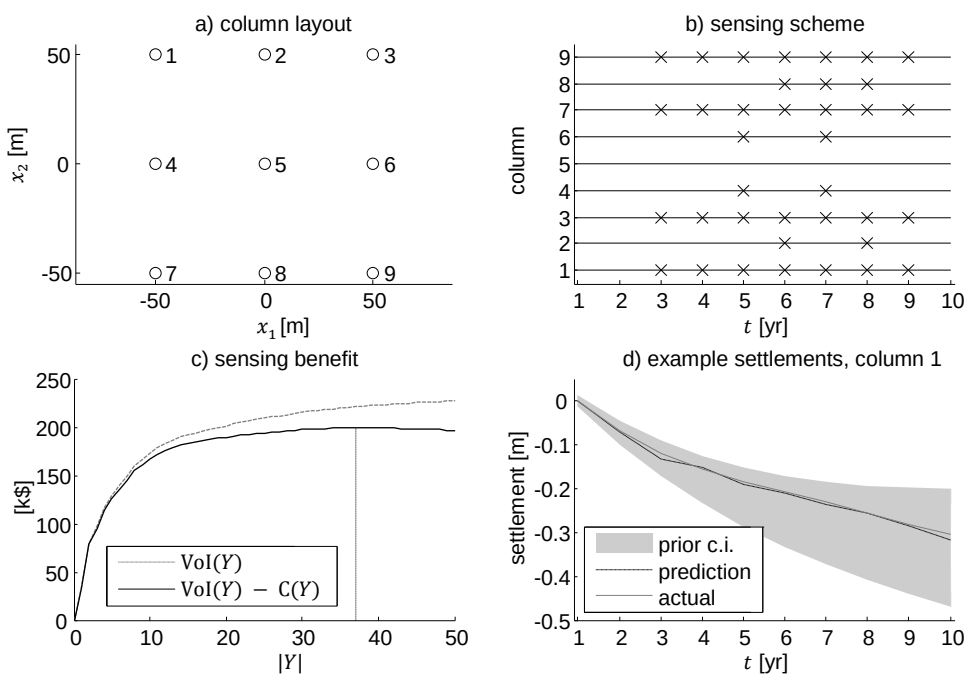

Figure 7-3: a) Layout of columns subjected to settlement; b) Optimal offline sensing scheme for this system, with measures denoted by x's; c) VoI for the scheme, with respect to the number of sensors greedily selected, with the optimal set indicated by the peak of the net VoI; d) Example settlement of column 1, with the prior 95% confidence interval (gray area), actual settlement (grey line), and settlement predicted by the optimal measurement scheme (dotted line).

The performance of online and offline sensing optimization are also compared for this problem. To evaluate online optimization, a specific instantiation of the settlement profile over time is used, as shown for column one in Figure 7-3d. Without any measurements, implementing the prior optimal management actions, a loss of $462k would be incurred by the managing agent following this particular realization of the settlement profile. With the offline optimal sensing

scheme of Figure 7-3b implemented, predicted settlements and management actions are altered based on collected data, reducing the loss to $276k. Following an online sensor placement and scheduling approach, where the optimal sensing plan is re-evaluated following the collection of settlement data in each year, the resulting loss is also $276k. However, the optimal online sensing scheme achieves the same loss reduction using 10 fewer measurements than the offline scheme.

Comparing across 20 different cases of online sensor placement, corresponding to 20 random instantiations of the system, in all cases the online optimal scheme consists of fewer measures than the offline optimal scheme, and in 15 cases the losses incurred by the managing agent are reduced as well. The adaptivity gap for this problem, i.e., the reduction in management losses and sensing costs for online sensing compared to offline sensing, ranges between 2% and 86% in these cases, with an average of 15%.

Finally, note that the sensing scheme is also sensitive to the type of data being collected. Consider an alternative formulation of this problem in which the measurements of settlement under each column have an unknown Gaussian bias. This bias can be modeled by introducing correlations between measurements taken on the same column at different times into the covariance matrix $\mathbf{\Sigma}_\epsilon$ of the measurement noise in Equation 3-6. In this example, an unknown systematic bias term with a Gaussian distribution with a mean of 0 and a standard deviation of 0.1m was introduced for the measurements associated with each column.

To eliminate such a bias, earlier measurements of settlement can be made when the settlement is known (from the prior model) to be low, such that this bias might be eliminated from later measurements. However, in this case the greedy optimization approach for selecting measurements performs relatively poorly. As shown in Figure 7-4a, the algorithm tends to cluster measures on relatively few columns. This is due to the fact that the first measure selected for a column has relatively little value due to the measurement bias. However, with this measure selected, the algorithm's next step is to select a measure for the same column at the first time step, so as to correct for this bias. Following this, the bias for all measures associated with this column can be removed, and so the value of these measures is higher, leading the algorithm to select all measures on this column before moving to another column and repeating the procedure. The ultimate consequence of the myopic greedy approach in this case is a sub-optimal set of sensor placements and schedules. This shortcoming in the greedy approach derives from the lack of submodularity in the VoI for this problem; the value of a pair of measures on a column (one at the first time step, and one later on) is higher than the sum of the values of each measure individually.

As discussed in Section 2.2.1, the standard greedy optimization approach employed throughout this work represents a forward greedy optimization approach, in which the optimal sensor set is built by iteratively adding measures to an empty set. However, in this case, a reverse greedy optimization approach, in which the full measurement set $\mathcal{Y}$ is iteratively pruned to remove elements, can outperform the forward greedy optimization approach. As shown in Figure 7-4a,

the set with the same number of sensors selected by the reverse approach allocates these sensors evenly across the corner columns of the structure. The reason for this is intuitive: by beginning with the complete set of measurements, the value provided by earlier measures in correcting the bias of later measures is captured, and these measures are preserved through the optimization process to retain this value. In Figure 7-4b, the net VoI provided by forward and reverse greedy optimization schemes is compared. For the case of 18 sensors, as shown in Figure 7-4a, the reverse approach outperforms the forward by 15%. Note, however, that the reverse greedy approach does not dominate the forward approach for all sensing set sizes in this case; both approaches remain suboptimal overall.

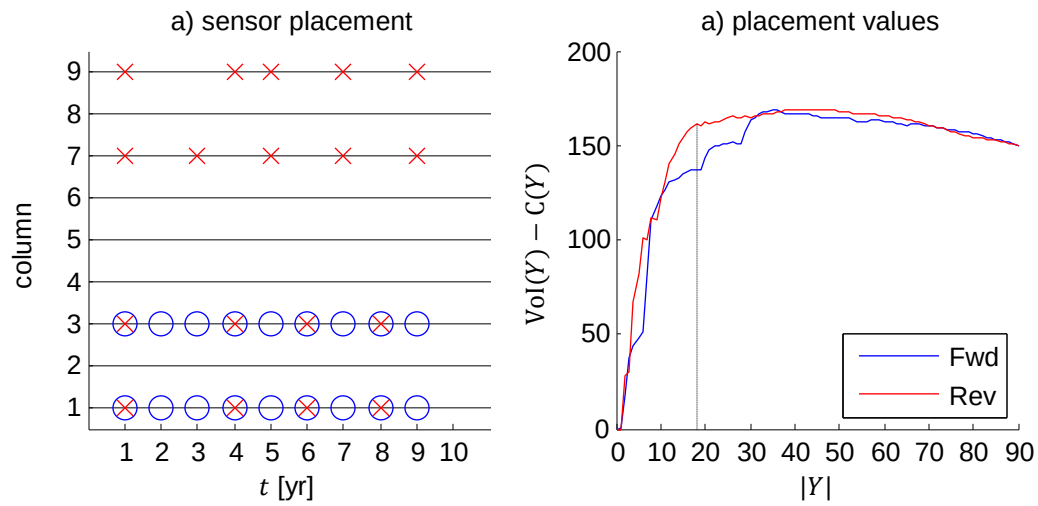

Figure 7-4: a) Comparison of forward greedy optimized (o) and reverse greedy optimized (x) sensor placement and scheduling schemes in the case of unknown measurement offsets; b) Comparison of the net VoI provided by sets of various seized optimized via forward and reverse greedy approaches; the vertical dotted line indicates the schemes depicted in (a).

The above discussion serves as a cautionary example of when the efficient but myopic forward greedy optimization approach leads to an overall suboptimal result for sensor placement and scheduling. As discussed, the forward approach is generally preferred for reasons of computational simplicity; beginning with an empty set and adding measures to arrive at an optimal set of a relatively small size is simpler than beginning with this full set of possible measures and iteratively removing many elements. In this particular case, however, the reverse greedy approach might be preferable, as it outperforms the forward approach, especially for relatively small numbers of sensors. The reason for the poor performance of the forward greedy approach can be readily understood in this example, and the benefits of the reverse approach likewise appreciated. However, in more complicated problems, such an intuitive understanding of the shortcomings of the optimization approach may be difficult. Furthermore, the reverse greedy optimization approach may not always represent a feasible or beneficial alternative. Further comparisons of these two greedy approaches are made in Section 11.1.

Chapter 8

Structural Health Monitoring of Scott Hall

This chapter demonstrates the application of the methodologies of optimal sensor placement developed over the previous chapters to a problem of SHM system design. The subject of this application is the Sherman and Joyce Bowie Scott Hall (or simply Scott Hall), a building on the campus of Carnegie Mellon University in Pittsburgh, Pennsylvania. Several major structural columns of this building were instrumented with fiber-optic strain gauges during the building's construction, and data collected from these sensors are used to demonstrate the methodology for optimal sensor placement using VoI.

It should be noted that this chapter presents an example of the optimization of sensor placements that makes use of data that are collected by sensors that are already in place at the time of the analysis. Therefore, this case study presents a "what-if" scenario, examining how these sensor placements might have been optimized if the probabilistic model of the structure developed from the collected data had been available prior to the placement of sensors. In principle, this probabilistic model might have been developed using computational models of the structural system, i.e., finite-element models. However, due to constraints of the project timetable and building construction, and lack of access to the design models for the Scott Hall building, no such analysis was conducted. Instead, this chapter can be interpreted either as an example of the application of the sensor placement methodologies discussed in this work to a structural system using a probabilistic model obtained from a "black box" source, or as an approach to the design of a more permanent sensing system using data collected by a cheaper but less durable exploratory sensing network. Finally, this caveat demonstrates one of the important practical barriers to sensor placement optimization, i.e., that the investment in time and collection of modeling information necessary to perform pre-posterior analysis may not be justified in some contexts, and that placements are instead often guided by engineering judgement as a "back of the envelope" suboptimal alternative.

8.1. Overview of Scott Hall and its Instrumentation

The Sherman and Joyce Bowie Scott Hall building began construction in the spring of 2012, and construction was completed by the spring of 2016, with the building officially opening on April 30, 2016. The building has about one hundred thousand square feet of floor space, and includes

classrooms and offices, collaborative spaces, a café, laboratories, and a ten thousand square foot cleanroom facility. Upon its completion, the building became the campus center for nanotechnology, biotechnology, and energy research.

The building itself consists of two structures; the "courtyard" component, containing the cleanroom and café, extends underground beneath the existing courtyard area between Wean, Hammerschlag, and Porter Halls. The second structure, which is of greater relevance to this chapter, is the "north wing", consisting of a concrete masonry structure supporting a four-story above-ground steel frame structure. This wing is constructed north of Hammerschlag Hall and west of Wean Hall, and part of the steel-frame structure extends out over a hillside, where it is supported by 12 circular steel columns, splaying out from 4 foundation pedestals. Figure 8-1 shows a computer rendering of the north wing structure, in which the tubular steel columns can be seen supporting the portion of the structure that is cantilevered over the hillside.

Figure 8-1: Computer rendering of Scott Hall's North Wing structure, between the existing Wean Hall (left) and Hammerschlag Hall (right). Image from http://www.cmu.edu/cdfd/scott-hall/.

Onto these 12 columns, 25 fiber-optic strain gauges were installed during the fall and winter of 2014. These fiber-optic strain gauges are produced by Micron Optics, Inc., and measure uniaxial strain along the direction of the fiber. An image of one of these strain gauges is shown in Figure

8-2a. These sensors operate on the fiber Bragg grating principle (Micron Optics, 2012). During the manufacture of the fiber optic strand, variations in the refractive index are introduced into the strand at regular intervals, forming the grating. Based on the spacing of these variations, a certain wavelength of light is more strongly reflected within the strand. As the strand is deformed, the spacing of these variations changes, and so the wavelength of the most strongly reflected light changes. A laser interrogator unit (shown in Figure 8-2b) sends pulses of light down a strand of fiber optic cable, and records the intensity of the light reflected at various wavelengths. By identifying peaks in the spectrum of the reflected light, the spacing of the variations in the strand can be computed, and this spacing can be compared to the original spacing of the grating to compute the relative strain in the strand. By using gratings with sufficiently different initial spacing, several gratings can be connected in series along a common strand, allowing for a single fiber-optic cable to support multiple sensors. The strain gauges used for this project (Micron Optics model os3155) also contain a second "free" grating, which is allowed to deform independently of the strain in the sensor (and therefore in the column to which the sensor is attached) to perform temperature compensation.

a) b)

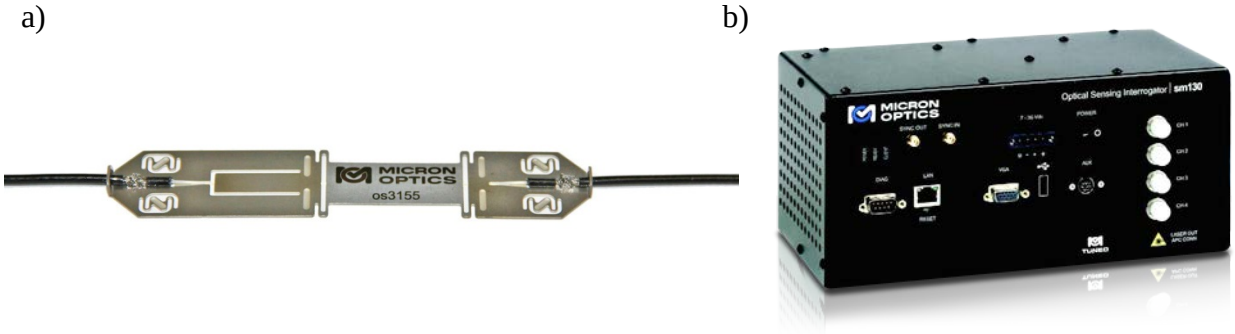

Figure 8-2: Micron Optics fiber-optic strain gauge (a) and interrogation unit (b). Images from www.micronoptics.com.

These sensors are attached via spot-welding to twelve circular steel columns supporting the cantilevered portion of the Scott Hall structure. Sensors are attached to the column ten feet above the point at which the columns are attached to their foundations via a pin connection, as shown in Figure 8-3. To either side of the sensor attachment location, metal tabs are welded to allow for the attachment of protective brackets. On all but one column, two sensors are attached to opposite sides of the steel column; this configuration allows for the assessment of axial strain in the column, as well as bending in the plane of the sensors. On one column (designated TC-E), three sensors are attached at 120º intervals around the column circumference; this configuration allows for assessment of the axial column strain as well as bending in two perpendicular directions.

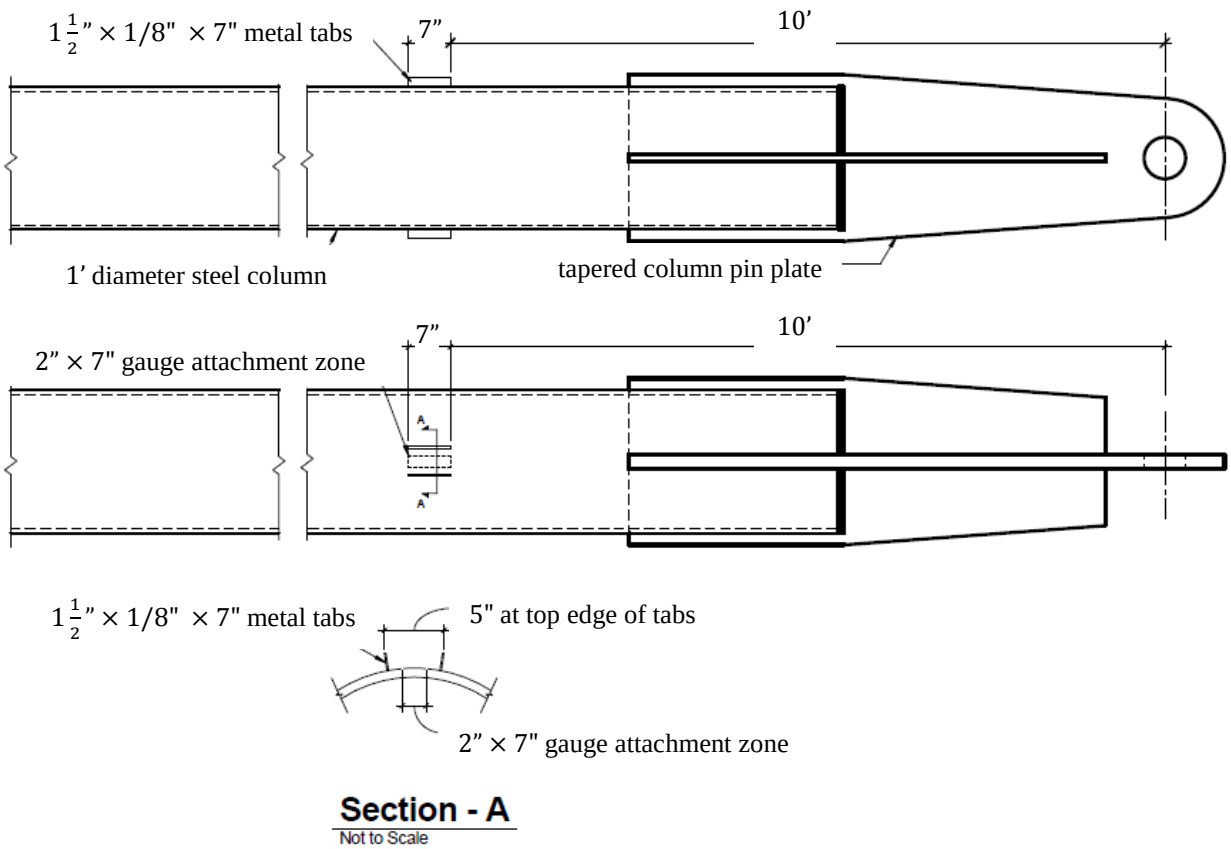

Figure 8-3: Diagram of strain gauge installation on structural columns.

Sensors were installed on the columns during November 2014. Images of the installed sensors are shown in Figure 8-4. Immediately following their installation, the sensors were covered with waterproof protective putty, which is visible as the green material in Figure 8-4b. Fiber-optic cables from each sensor were connected to those of neighboring sensors on the column, and to longer cables to reach down to the base of each column. Shortly after installation, the sensors were connected to the interrogation unit for testing, and to obtain a baseline strain measurement for the installed sensors, against which subsequent strain readings could be compared to detect changes in column strain.

a) b)

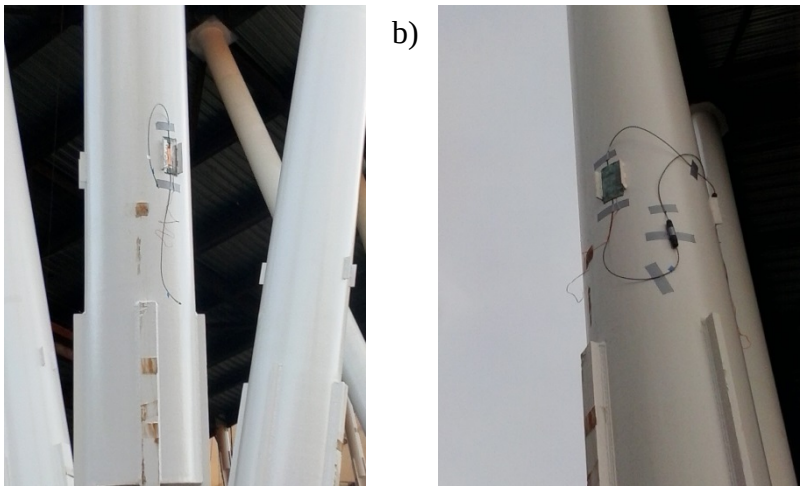

Figure 8-4: Photographs of installed sensors on column TC-E from the northwest (a) and southwest (b).

Once all sensors had been installed, a temporary configuration was established to collect data from the sensors during the construction of the structure. Due to the need to keep construction areas clear, this configuration was disassembled intermittently and re-assembled when possible. A diagram of the final sensor configuration, indicting how sensors are connected to each other and to the interrogation unit (I.U.) is shown in Figure 8-5. During the construction period, the temporary configuration also followed this scheme, except that the two sensors on column TC-M were not connected, due to insufficient fiber optic cable being available to join these sensors into the network. Therefore, during the construction period, only data from 23 sensors on eleven columns were collected.

Strain measurements from these sensors were collected during the construction period at a frequency of 1 Hz over two extended periods; from November 26 to November 30 (for a total of about 5 days of data) and from December 8 to December 15 (about 7 days). It is these two datasets that are used in the remainder of this chapter to model the column strains and to validate the proposed methodology for optimal sensor placement.

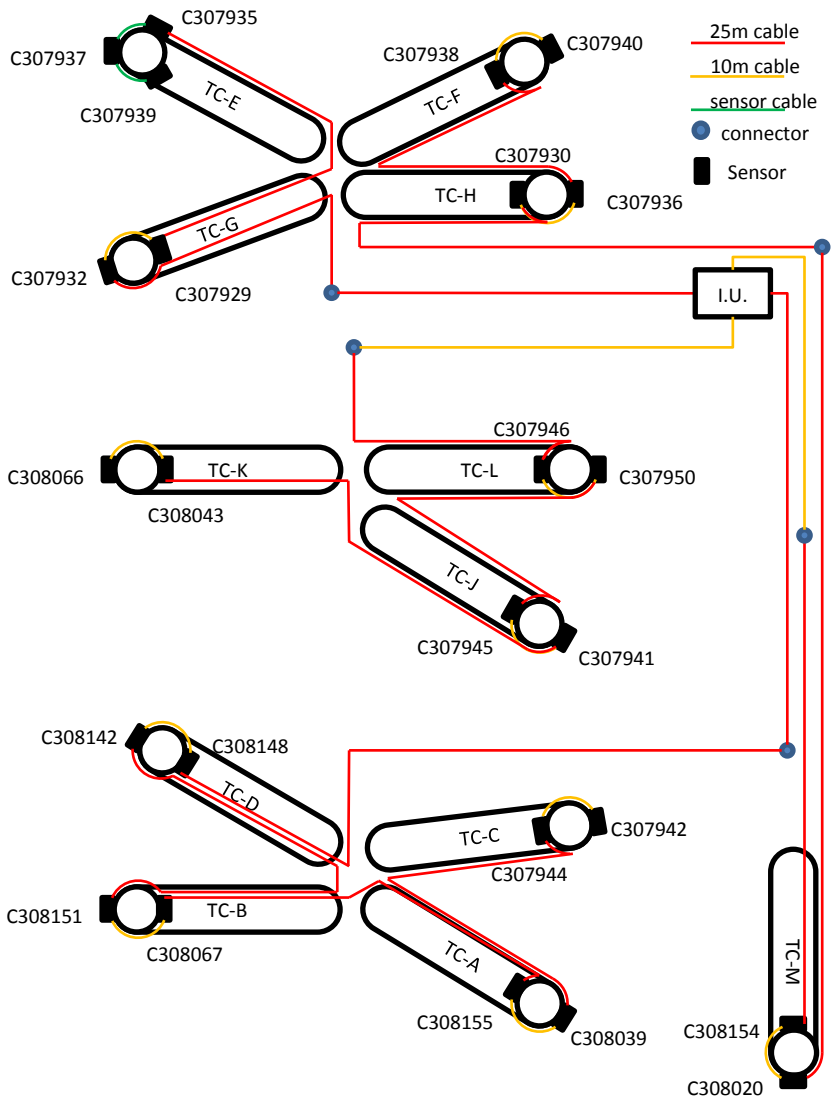

Figure 8-5: Diagram of sensors, columns, and fiber optic cable configuration.

8.2. Probabilistic Strain Field Modeling

Using data collected by sensors during construction, a probabilistic model of the strain field in the columns is created. Rather than modeling the strains themselves, the linear strain rates in the columns for one-hour periods are modeled as Gaussian random variables. This is done to support a decision-making problem concerning the relative rate of strain in the columns, and the result of these rates on the deformation of the building deck. During the construction phase, such differential strains occur as new dead loads are added unevenly to different parts of the structure. In a finished structure, such trends might indicate differential settlement of the foundations; columns supported by a settling foundation would experience decreasing strain relative to those

supported by other foundations. Therefore, the problem of addressing differential strain and loading is relevant both during and following the construction phase.

The underlying random field variables $\mathbf{f}$ are defined as the strain rate $\mathbf{f}_\theta$ (in micro-strain per hour) and starting strain $\mathbf{f}_\phi$ (in micro-strain) for all columns for any one-hour period, assuming that column strains follow a linear trend during this period. These variables are described by a joint Gaussian distribution:

$$\mathbf{f} = \begin{bmatrix} \mathbf{f}_\theta \\ \mathbf{f}_\phi \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \boldsymbol{\mu}_\theta \\ \boldsymbol{\mu}_\phi \end{bmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_\theta & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}_\phi \end{bmatrix} \right) \tag{8-1}$$

The parameters of this distribution are obtained by analyzing the strain data collected during the first five-day data collection period. These data are divided into one-hour segments, the strain rates and initial strains for each column during each hour are computed, and these are used to identify the mean vectors and covariance matrices for these variables, assuming a Gaussian distribution. Also note that it is assumed that the strain rates and starting strains are independent.

The strain measurements collected by a given subset of the 23 strain sensors for which data were collected during construction are used to update the column strain rate model. That is, for a given hour of time, the data collected by sensors during this hour are used to define a posterior distribution over the column strain rates and initial strain values. Measurements are defined as in Equation 3-5, with each row of $\boldsymbol{\Omega}_Y$ encoding the timestamp of the corresponding measurement in an appropriate position (such that it will multiply with the correct entry of $\mathbf{f}_\theta$) and having a one in an appropriate position (such that it will multiply with the correct entry of $\mathbf{f}_\phi$). Measurement noise is assumed to be independent between different sensors and at different times, and the noise level is set for each sensor based on the difference of strain measurements from the assumed linear trend strain model. In this way, the measurement vector $\mathbf{y}$ concatenates the data collected by the selected subset of 23 sensors over any one-hour period. Representative measurements $\mathbf{y}$ are obtained by dividing the second seven-day data collection period into one-hour segments and using the measurements of the selected sensors during each of these segments as a sample of $\mathbf{y}$, for use in the sampling-based evaluation of the expectation as in Equation 5-25.

Limit-state variables are defined following Equation 3-11, and describe relative displacement between pairs of representative locations on the structural deck. Each row of $\boldsymbol{\Xi}_G$ encodes the relationship between the column strains and the relative vertical displacement between a pair of representative deck locations. Three representative locations are considered at the south, center, and north ends of the structural deck, and therefore three relative displacements between pairs of locations are considered, i.e., the limit-state variables define $n = 3$ relative displacement rates between these representative deck locations. These relationships are determined from a simple linear-elastic finite element model of the building, which includes the eleven measured steel columns and the deck supported by them.

Component states in this model correspond to whether or not the differential displacement rates of the representative structural deck locations are within a specified bound (arbitrarily designated as 4 mm/hr for this problem). State variables are therefore defined as in Equation 7-21, with appropriate upper- and lower-bounds for each limit-state variable (i.e., for each differential displacement rate). The state of the system is then defined using a series system of the components, i.e., if the differential displacement rate of any two representative points exceeds its threshold then a failure is assumed to occur.

Figure 8-6 shows an example of the prior and posterior distributions for one of the limit-state variables in this model, the differential displacement between two representative structural deck locations. The prior distribution is centered about zero, and there is about a 2% prior probability of this rate exceeding its bounds for any given hour. However, using a sample of the collected data, this distribution is updated to a posterior with a mean of 7 mm/hr, well above the threshold. While the true trend is about 8 mm/hr, the posterior model correctly predicts that the probability of differential loading causing a differential displacement of the structure is close to 100%. The data used for this example were collected during a time interval when concrete was being poured in one area of the structure; therefore a differential loading was in fact occurring at this time, resulting in the observed differential displacement trend.

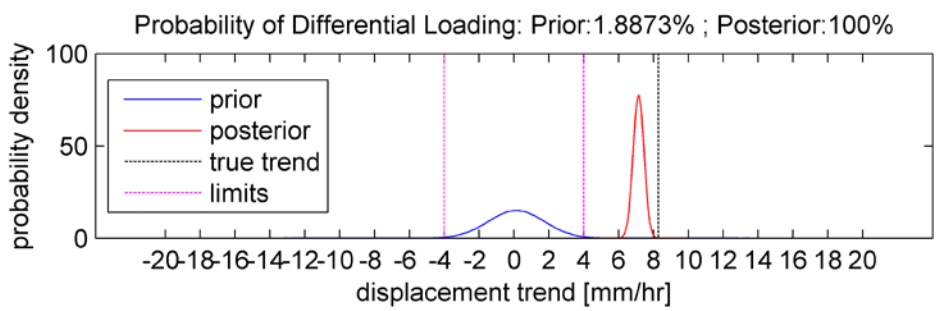

Figure 8-6: Example of sensor data updating the prior distribution over relative displacement trends for a time interval to a posterior distribution. The true trend for the time interval computed from the data is indicated for reference.

A decision-making problem is defined based on this differential displacement model. The agent managing the structure has the option to halt construction, and thereby prevent differential loading, at a cost of $1M. Alternatively, if construction is not halted and differential displacement above the specified limit is experienced, a failure cost of $10M is incurred. The loss function of Equation 4-16 is therefore used in this problem, with costs as stated. Note that these values are merely representative numbers used to define the decision-making problem for this example; no such costs were incurred during the construction of Scott Hall.

8.3. Sensor Placement Results

First, to address the optimal placement of sensors, the management problem is considered to be purely spatial. That is, the decision-making problem for a representative hour-long period is considered and used to define the VoI metric to guide the selection of a subset of the 23 sensors active during construction that would have been optimal to support the given decision-making problem. This treatment is consistent with the assumption that, at each hour, an independent sample of **f** is realized from its prior distribution, data are gathered, and a management decision is made. Sensor placements can be optimized to support decision-making during any typical hour-long period. In this way, a spatio-temporal system (the actual structure, whose strains vary both in space between columns and through time as construction proceeds) is treated as a purely spatial system, and sensor placement is performed for any representative time step. The results when the system is treated as spatio-temporal are discussed in Section 8.4.

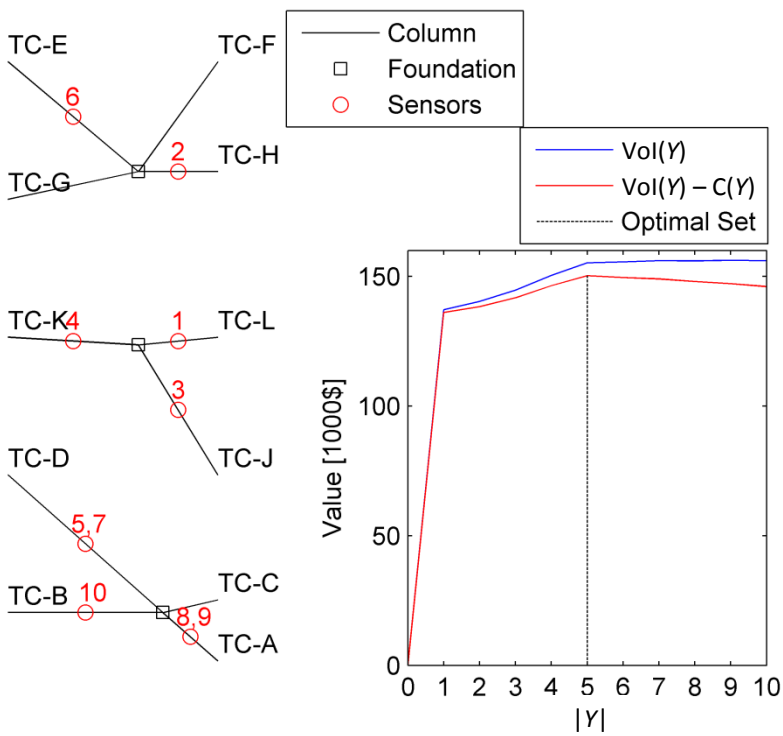

Figure 8-7: Optimized sensor placements (left) and associated VoI (right).

Figure 8-7 shows the results of optimal sensor placement using the VoI metric based on the decision-making problem outlined above. To the left, the locations of the optimal sensors and the columns to which they are attached are indicated, along with the order in which the sensors are greedily selected. To the right, the VoI provided by these greedily selected sets are indicated, as well as the net VoI assuming a cost of $1k per sensor. Under this cost, the optimal set consists of 5 sensors on tubular columns L, H, J, K, and D. This set consists of one column associated with each of the three common foundations. Furthermore, the sensors selected for the optimal set tend

to be on columns that are most sensitive to differential strain trends, as determined from the linear-elastic finite-element model. For example, the first and second sensors selected are on columns TC-L and TC-H, whose strains are the most sensitive to the displacements of their respective foundations (as they are the most nearly vertical of the columns associated with these foundations). However, the measures collected by these sensors are also correlated with those obtained by sensors on other columns as well. It is therefore a tradeoff between sensitivity and shared information that leads to the optimal sensor configuration by this objective.

8.4. Sensor Placement and Scheduling Results

Next, this problem is treated from a spatio-temporal perspective, as in Chapter 7. A spatio-temporal multivariate Gaussian model is developed for $\mathbf{f}$, where the strain rate and initial strain for each hour over a 24 hour period are associated with a separate temporal instantiation of the variables, with correlation between the variables of different time steps, as observed in the five-day data collection period. Strain measures of the columns are again obtained for one-hour periods. In this case, the choice of sensor placements corresponds to the selection of a subset of the 23 active sensors to monitor, while that of sensor scheduling corresponds to selecting which hour of the day to collect measures. Decision-making for each hour-long time step is defined as before, but with costs reduced to \$10k per hour and \$1k per hour for failure and intervention, respectively. Over the 24-hour period, the total cost is assumed to be the sum of costs in each hour, defining a temporally decomposable loss function as in Equation 7-4. Note that due to the series topology, the loss does not decompose spatially. No discounting is considered, i.e., $\gamma(t) = 1$. Decision-making is performed using only information gathered up to and including the previous hour, i.e., data collected during a given hour cannot support the decision-making in that hour.

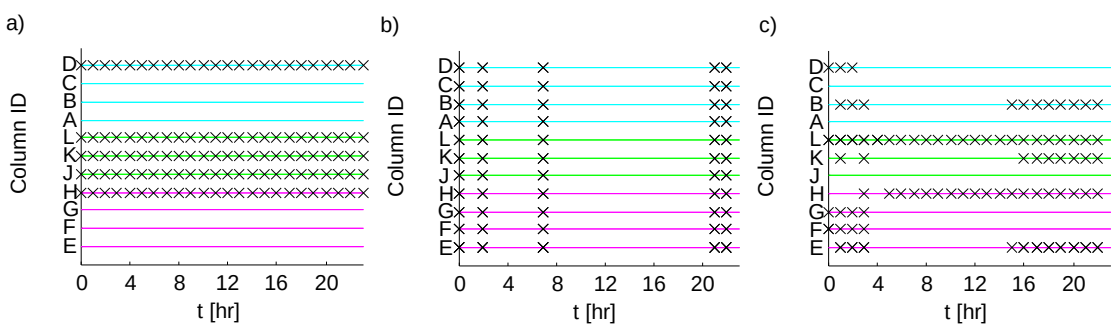

Figure 8-8: Results for sensor placement (a), scheduling (b), and placement and scheduling (c).

For this spatio-temporal decision-making problem, three sensing schemes are optimized for a typical day: an optimal placement, optimal schedule, and optimal placement and schedule. Results for each of these are shown in Figure 8-8. In Figure 8-8a, the optimal placement scheme, indicating which columns should be monitored throughout the day, is obtained by constructing

the sensing cost function such that measurements by a previously unselected sensor have a positive cost ($1k, as before), but measurements by previously selected sensors at a different time of day have no additional cost. It should be noted that this optimal placement exactly matches that of Figure 8-7, with columns D, L, K, J, and H being monitored. This is to be expected, since this scheme was optimized for any typical one-hour period, and therefore it is optimal for any specific time step in this spatio-temporal case.

Figure 8-8b depicts the optimal sensing schedule, i.e., the times during the 24-hour period when it is best to collect measures from all sensors. In this case, the sensing cost function is constructed such that the selection of a measure at a different time from previously selected measures has a positive cost, but the selection of a new measure at the same time has no additional cost. Here, some measurements are taken earlier in the cycle so that data will be available to guide later decision-making. Also, these measures correspond to the times of day (late morning and early afternoon) when construction activity is most intense, and therefore the probability of differential loading is highest (note that the 24 hours of the planning time period do not correspond to the hours of the day, but instead to the time after the start of the cycle, which is at noon).

Figure 8-8c depicts the optimal sensor placement and scheduling scheme. This scheme was optimized using a cost function that puts a price on each new measure, regardless of whether previous measures on columns or other measures at the same time have been made. Here, prescribed measurements switch between columns as time passes, adjusting for changes in loading and strain patterns over time. Note that, as in the optimal placement, columns L and H tend to be frequently monitored, while as in the optimal schedule, measures are clustered at the beginning and end of the cycle, corresponding to the times of greatest on-site activity.

The VoI of the sensor scheduling scheme is $1.6k, that of the placement scheme is $2.8k, and that of the optimal placement and scheduling is highest at $3.2k. While all schemes prescribe roughly the same number of sensor measurements, the optimal sensor placement and scheduling is the most flexible (as the cost structure does not incentivize the inclusion of all measures of a particular sensor or at a particular time), and therefore can provide greater benefits in terms of reduced management costs (as measured by VoI) while using a comparable number of measures. Note that the VoI for the optimal placement differs from that of Section 8.3, due to the different cost structure.

Finally, note that the schemes prescribed here are obtained using a probabilistic model that is trained from only a five-day-long data record, i.e., the training period only includes five samples of the daily strain cycle. These results may therefore be biased by this small training set, focusing on appropriately responding to loading anomalies that occurred during this period, rather than providing a more generally optimal scheme. This is a concern for all situations where models must be trained using limited data.

Chapter 9

Seismic Risk in the San Francisco Bay Area

This chapter presents an application of the optimal sensor placement methodology to the monitoring of seismic risk, supporting the management of an infrastructure system of bridges and tunnels in the San Francisco Bay area. The hazards associated with earthquakes, in terms of the spatial distribution of ground motion resulting from events with known epicenter locations and magnitudes, are well-studied phenomena, both in general and for the San Francisco Bay area in particular. Existing probabilistic models, including Poisson process models for earthquake recurrence (e.g., Anagnos and Kiremidjian, 1988) and ground motion attenuation equations (e.g., Douglas, 2011), are therefore used here to define the spatial distribution of the demand placed on an infrastructure system due to a seismic event. Furthermore, the response of infrastructure components to these ground motions have been extensively studied, and probabilistic models for this response, such as fragility curves commonly used for seismic damage estimation (e.g., Hazus-MH, 2013), are used to define probabilistic models for the capacities of infrastructure components to resist seismic demand. Based on these models, and a decision-making problem concerning the closure of bridges and tunnels in the aftermath of a seismic event, an optimal arrangement of sensors for supporting system management is proposed. These sensors include both demand measurements, i.e., seismic monitoring at the locations of infrastructure components to better determine the demands placed on these, and capacity measurements, i.e., improved assessments of the components in question (e.g., through more regular and thorough inspections or the installation of SHM systems) to better determine their capacity with respect to seismic loadings.

This example is meant to illustrate the application of the sensor placement methodologies discussed on a larger scale, dealing with multiple separate structural systems as components distributed across several kilometers within the spatial domain. Furthermore, rather than using a full spatio-temporal model of the system as in Chapter 7, a purely spatial problem is considered, with the temporal effects of seismic risk being easily treated within the context of the Poisson process earthquake recurrence model and some basic assumptions on the decision-making problem, as will be discussed in Section 9.1.1. Additionally, in Section 9.2.1 a VoI density analysis is presented, which provides additional insight into how particular seismic events, based on their magnitude and epicenter locations, contribute to the value of the optimal sensor

configuration. These insights can lead to future efforts to more efficiently sample scenarios for the estimation of VoI via Monte-Carlo methods.

It should be noted that this example is presented for illustrative purposes only. Although the system considered and its constituent components are based on real-life infrastructures, and applicable seismic hazard and structural response models have been used, the results are not intended as a practical recommendation for the management of these infrastructures. The decision-making problem discussed here, including the costs of management options and the costs of sensing efforts, does not necessarily reflect the true actions and costs associated with the management of these infrastructures. Furthermore, network effects, i.e., the interdependent effects of the closures of components of the transportation system on the capabilities of the system as a whole, are ignored in this problem to allow for efficient VoI computation using the decomposable loss assumption, as discussed in Section 5.2. For these reasons, a more complete and focused analysis should be performed before any of these recommendations can be applied to support management of infrastructures in the San Francisco Bay area.

9.1. Seismic Risk Model

This section outlines the probabilistic model used to describe seismic risk in the San Francisco Bay area and the responses of structures, as well as the monitoring of these variables and the decision-making problem defined for this system. The basic characteristics of the seismic risk are defined using model parameter and scenario variables W, as introduced in Section 2.1.1.1. A seismic scenario, i.e., a specific earthquake event, is parameterized (for the purposes of this application) by its magnitude, denoted w_m (a scalar quantity), and epicenter location, denoted $\mathbf{w}_e$ (a vector of coordinates). Earthquakes affecting multiple faults and/or those with distributed epicenters are not considered. Both of these are treated as random variables defining the scenario, i.e., $\mathbf{w} = \{w_m, \mathbf{w}_e\}$. The prior distribution p_W over these scenario parameters is a homogeneous Poisson process model, based on the model outlined by Anagnos and Kiremidjian (1988) and using information for the San Francisco Bay area as provided by Field et al. (2009). First, this model assumes independence between the magnitude and epicenter variables, i.e., $\mathrm{p}_W(\mathbf{w}) = \mathrm{p}_{W_m}(w_m)\mathrm{p}_{W_e}(\mathbf{w}_e)$. The distribution over the epicenter locations p_{W_e} is defined using a uniform distribution over locations along the main fault lines in the San Francisco Bay area, weighted by the relative probabilities of earthquakes originating from each fault, as in Field et al. (2009). The distribution over magnitudes p_{W_m} is defined using the recurrence rates of earthquakes of different magnitudes. That is, an earthquake with magnitude w_m or greater will occur in the regional at an annual rate $\Lambda(w_m)$. For this example, a lower bound of $w_m = 3$ and an upper bound of $w_m = 8.4$ are used, following Field et al. (2009). These recurrence rates can be used to define a truncated exponential distribution for the probability of an earthquake having a specific magnitude within this range, following the Gutenberg-Richter law (Gutenberg and Richter, 1954). Finally, as discussed in Section 2.1.1.1, these parameters are assumed to be known to any decision-maker without the need to infer them through sensing. Practically, this

assumption is reasonable, as earthquake information is determined with high accuracy in the aftermath of an event by existing monitoring systems operated by the USGS, and this information would then be available to the decision-makers.

Conditional to the earthquake magnitude and epicenter location, the impact of this event on infrastructure is defined following Equation 2-2. This impact is quantified through two sets of random variables: the demands placed on the components of the system by the seismic event, denoted as **d**, and the capacities of these components with respect to this demand, denoted as **c**. Demands are defined in terms of the peak ground acceleration at the component locations, using ground motion attenuation equations developed by Boore and Atkinson, as presented by Douglas (2011). In this model, ground accelerations are log-normal random variables whose parameters are functions of the distance to the epicenter location and the earthquake's magnitude, as well as region-specific soil properties. However, as discussed in Section 3.3, under the logarithm transformation these accelerations can be described using a multivariate Gaussian distribution. Correlations between demands at different locations are assumed to be given by an isotropic squared-exponential covariance function, as in Equation 6-3, with a length scale parameter of $\lambda = 10$ km. Similar models for correlation of earthquake acceleration are used by, e.g., Bensi et al. (2015). Although more advanced models for spectral acceleration correlations exist (e.g., Loth and Baker, 2013), this simplified model is applied for illustrative purposes. Seismic capacities are also modeled as log-normal random variables, using the log-normal fragility curves defined for various types of infrastructure components in the Hazus-MH software (2013). Correlations between capacity variables for different components are assumed to be higher for components with a similar overall typology, i.e., components that share similar structural characteristics, materials, or designs. Recent work has begun to investigate these capacity correlations for reinforced concrete structures under seismic risk using dynamic structural models (Zelaschi et al., 2017). This work provides support for the idea that bridge capacities are correlated between components of a similar typology, although identifying an appropriate model for this correlation is still an open area of research. Therefore, in this example, values of the correlation coefficients between various component types are set using engineering judgement, with the constraint that the resulting covariance matrix is positive semi-definite. These assumed correlation coefficients are listed in Appendix C for reference.

In all, $n = 27$ components are considered in this example, corresponding to major bridges and tunnels in the San Francisco Bay area (these are listed in Appendix C for reference). As discussed, the random field variables consist of the log-capacity and log-demand variables, and so the vector of these variables is:

$$\mathbf{f} = \begin{bmatrix} \log(\mathbf{c}) \\ \log(\mathbf{d}) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \boldsymbol{\mu}_c \\ \boldsymbol{\mu}_{d|\mathbf{w}} \end{bmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_c & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}_{d|\mathbf{w}} \end{bmatrix} \right) \tag{9-1}$$

Note that the mean and covariance of the log-demand variables are functions of the seismic scenario parameters, while those for the log-capacity are not. Limit-state variables are defined

for each component following Equation 3-11, with $\mathbf{\Xi}_G = [\mathbf{I} \quad -\mathbf{I}]$ and $\mathbf{g}_0 = \mathbf{0}$, i.e., the limit-state variables are the differences between the log-capacity and log-demand for that component. Likewise, component failures are considered to occur whenever demand exceeds capacity, with binary state variables being defined for each component as in Equation 3-14. Note that a failure state for a component refers not necessarily to an immediate structural collapse (which might be readily observed without the need for sensors) but to severe damage to critical elements of the structure, which would lead to such a collapse if the structure were kept in service. That is, a "failure" refers in this case to the structural integrity of a bridge or tunnel being severely compromised, but in a way that is not readily apparent to casual observation.

Measurements of the capacities of each component and/or the ground accelerations at their locations are assumed to be possible, with measurements of the log-transformed quantities defined as in Equation 3-5. That is, measurements are assumed to have multiplicative log-normal noise (with median value 1 and coefficient of variation 0.2, independent between measurements) such that the transformed measurement noise is additive and Gaussian (with zero mean). Potential measurements are also assigned a cost, consisting of the installation and ongoing monitoring and maintenance costs for the sensors in question. It is assumed that measurements of a component's capacity, which might correspond to frequent in-depth inspections or the deployment of an SHM system, have a higher cost of \$46M over the lifetime of the sensors (which represents a recurring annual cost discounted to present value at a 5% rate) than measurements of the demand, which are estimated at \$23M over the lifetime of the sensors. A fixed sensing budget of $b = \$200\text{M}$ is used to constrain the optimization of Equation 2-16.

To define a decision-making problem for the management of this system, it is assumed that binary decisions must be made for each component in the aftermath of an earthquake as to whether or not to close each bridge or tunnel to traffic. The option to close the ith component incurs a certain closure or response cost $C_{r,i}$ for the managing agent, corresponding to the service level reduction for the transportation system, losses in toll revenues, etc., until appropriate full-scale inspections and repair efforts can be implemented on the structure. These values are roughly estimated for each component considered, with values listed in Appendix C. Note that network effects of the service disruption are not considered. If a component is not closed, no cost will be incurred unless the earthquake compromised the component's integrity, in which case a high failure cost $C_{f,i}$ is incurred for the ith component (costs are again listed in Appendix C). This problem defines a component-level loss as in Equation 4-15, and therefore a system-level loss that is decomposable across components as in Equation 3-1, allowing for efficient evaluation of the VoI as discussed in Section 5.2.

9.1.1. Value of Information in Poisson Processes

Earthquakes represent extreme events that recur at uncertain intervals. Although seismic risk is present at all times (and is assumed to be stationary in time for this example), earthquakes themselves are discrete events, which recur at uncertain time intervals. Furthermore, each event

is different, with a different earthquake epicenter and magnitude defining a different demand field, as discussed previously. Therefore, rather than optimizing sensor placements for a specific event, it is necessary to optimize these placements under uncertainty in the parameters of the event. Furthermore, rather than computing the VoI of sensor placements for supporting the response to any single earthquake event, this metric should be computed taking into account that these events will recur over time. Fortunately, rather than needing to resort to a full spatio-temporal model of the risk and system response as in Chapter 7, the properties of the Poisson process model adopted here allows VoI to be computed for this recurring seismic risk in a straightforward manner.

First, to optimize sensor placement under uncertainty in the parameters (i.e., magnitude and epicenter) of the scenario that will occur, the VoI for a representative or typical event should be evaluated as the expected value of the VoI computed under a specific event parameterized by $\mathbf{w}$, denoted $\mathrm{VoI}(Y|\mathbf{w})$, across possible values for $\mathbf{w}$ described by the prior distribution p_W. That is, the VoI provided by measurement scheme Y for a typical event is:

$$\mathrm{VoI}_{event}(Y) = \mathbb{E}_W \mathrm{VoI}(Y|\mathbf{w}) \tag{9-2}$$

This represents one approach to defining a VoI metric that is robust under uncertainty in W. An alternative approach is to define the VoI using a worst-case scenario, i.e., to replace the expecation over W with a minimization. However, it may be the case that, for each proposed sensor placement scheme, a scenario exists where there is no VoI for that scheme, and therefore no scheme could optimize this objective. For this reason, the expected value formulation shown above is adopted, and a minimization-based alternative formulation is not investigated.

In practice, the expectation of Equation 9-2 is approximated by drawing a number n''_{sim} of samples of the parameters from prior distribution p_W and estimating the VoI of a proposed sensing scheme Y using these samples as:

$$\mathrm{VoI}_{event}(Y) \approx \frac{1}{n''_{sim}} \sum_{j=1}^{n''_{sim}} \mathrm{VoI}\left(Y|\mathbf{w}^{(j)}\right) \quad \left\{\mathbf{w}^{(j)}\right\}_{j=1}^{n''_{sim}} \sim \mathrm{p}_W \tag{9-3}$$

where $n''_{sim} = 1000$ in this example. This number of samples represents a trade-off between stability of the approximation and computational speed. The corresponding coefficient of variation in the estimates of the VoI was empirically determined to be about 5%. The variance of this approximation can have an impact on the optimal sensor placement scheme selected since, if this variance is high with respect to the VoI, it will no longer be clear which sensor placement scheme is optimal. However, since the same set of scenario samples is used to compute the VoI for each proposed sensor placement scheme, this impact is minimal.

Second, in order to account for the recurrance of the risk event, the single event VoI above must be transformed into a discounted cumulative VoI for the measurement scheme over the lifetime of the system. Under the Poisson process model assumptions used here, this can be done by

using a simple multiplying factor, following Takahashi et al. (2004) and Pozzi and Der Kiureghian (2011a):

$$\mathrm{VoI}_{discounted\ cumulative}(Y) = \frac{\Lambda_{annual}}{1-\gamma_{annual}} \mathrm{VoI}_{event}(Y) \tag{9-4}$$

where Λ_{annual} is the basic Poisson process rate (in this case, $\Lambda_{annual} = \Lambda(3)$, i.e., the annual recurrance rate of earthquakes magnitude 3 and greater) and γ_{annual} is the annual discounting factor used to evaluate the present value of future costs (in this case, a 5% rate is used, i.e., $\gamma_{annual} = 0.95$). This formulation can be understood by noting that, in an expected sense, the VoI for a single event is realized whenever an event occurs, and therefore the expected VoI per year is the product of this single event VoI and the expected number of events per year. This discounted cumulative VoI is the benefit that is traded off against the lifetime sensing cost (again discounted to its present value) in determining the optimal sensor set, as in Eqution 4-12.

The formulation of Equation 9-4 relies on several assumptions on the nature of the seismic risk and the decision-making conducted to manage the system in response to this risk. First, this formulation assumes an infinite time horizon, although this might be used to reasonably approximate the VoI for finite horizons, especially if the discount rate for future events is relatively high. Second, this assumes that the system is unchanging over time, i.e., that the same probabilistic model will describe each event and the responses of the components, the same sensors are in place for monitoring the system, and the same management actions will be available with the same costs. In practice, the configuration of an infrastructure system, the risks to which it is subjected, and the ways in which it is monitored and managed will change over time. Furthermore, if components are damaged in one event, they may be repaired to different standards or entirely replaced before the next event, changing the characteristics of the system. However, with an appropriate discount rate, benefits accrued in the near-future when the system and its management remain roughly the same can be emphasized. Third, in a related point, independence between successive events is assumed. However, in practice, phenomena such as aftershocks or the possibility of multi-fault ruptures alter the probabilistic character of the risk in the aftermath of an earthquake. Furthermore, since the same components are subjected to successive events, their responses will be coupled. Although the assumption of complete independence between subsequent events is questionable, the benefits in terms of simple evaluation of life-cycle VoI of a monitoring scheme are assumed to be sufficient to justify this approximation, and are therefore employed in this example application.

9.2. Application to San Francisco Bay Area

Based on the infrastructure system and decision-making problem described above, optimal sensor placement is performed. In addition to evaluating the VoI metric for this problem, the conditional entropy metric, using the sum of marginal component state variable entropies as in Section 4.4.3.4, is used. This is done in order to provide a point of comparison for the placements

via the VoI metric, and also to compare the performance of these alternative metrics in this problem.

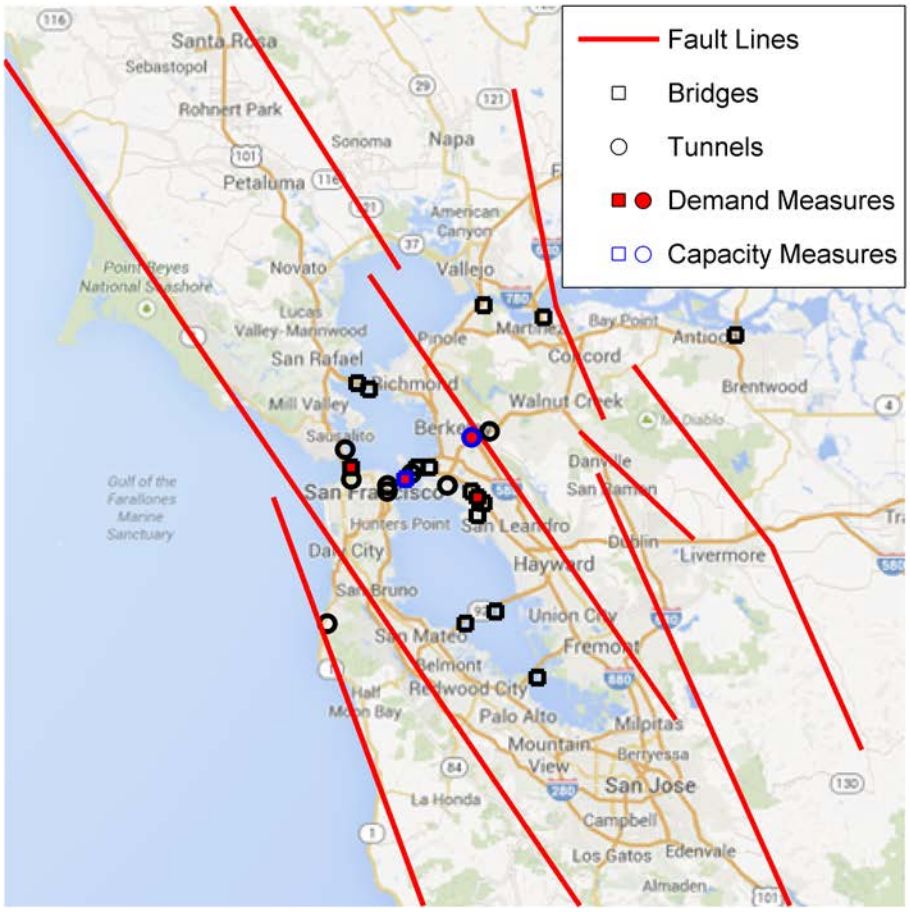

Figure 9-1: Proposed sensor placements based on the conditional entropy metric. Background image from www.maps.google.com.

Figure 9-1 shows the measurements selected by greedy optimization using the conditional entropy metric, which are listed for reference in Table 9-A, based on their order of selection. Six monitoring systems are placed (two for measuring capacity and four for measuring demand) before the allocated budget is exhausted. All demand measurements are in the upper San Francisco Bay area, where the concentration of infrastructure components considered is highest. This allows these demand measurements to reduce the posterior uncertainty in the demand for many of these nearby components, and therefore reduce posterior uncertainty in their states. Capacity measurements are prescribed for one bridge (the western span of the Oakland Bay Bridge, a steel suspension bridge) and one tunnel (the Northbrae Tunnel, a cut-and-cover tunnel). The capacities of these components are correlated with those of many other similar bridges (most bridges considered are steel) and tunnels (most tunnels considered are cut-and-cover), and so measurements of these components' capacities will contribute to reducing the posterior

uncertainties in the capacity variables of many components, and therefore in reducing the uncertainties in their states. Qualitatively, this is consistent with the conditional entropy metric's general goal of reducing uncertainty.

Table 9-A: Sensor placements by the conditional entropy metric, in order of greedy selection.

	Component Name	**Measure**
1	Golden Gate Bridge	Demand
2	Northbrae Tunnel	Capacity
3	Northbrae Tunnel	Demand
4	Park Street Bridge	Demand
5	San Francisco Oakland Bay Bridge (West Span)	Capacity
6	San Francisco Oakland Bay Bridge (West Span)	Demand

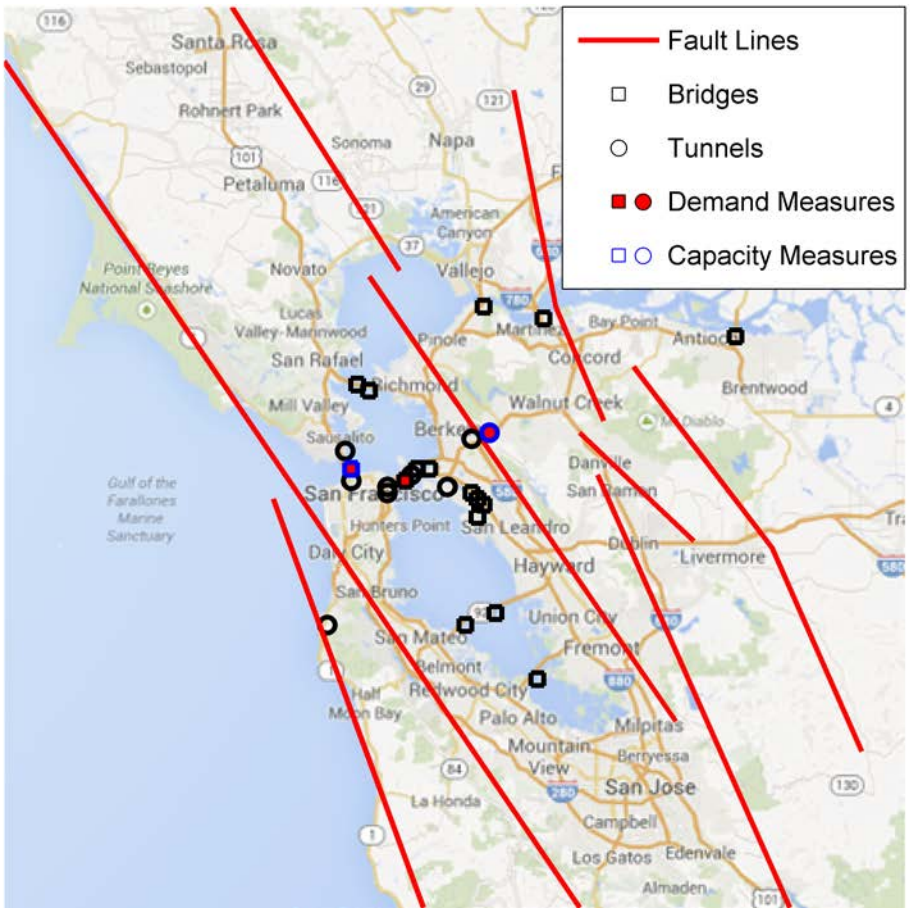

Figure 9-2: Proposed sensor placements based on the VoI metric. Background image from www.maps.google.com.

Figure 9-2 illustrates the sensor placements obtained after optimization using the VoI metric, with the selected measurements listed in Table 9-B. Here, five measurements are recommended,

two to measure capacity and three to measure demand. Rather that meeting the budgetary constraint, no additional sensors are placed because, beyond this point, the benefits of adding sensors (as measured by the VoI metric) are less than the costs of the sensors.

Demand measurements are again concentrated in the upper San Francisco Bay area, but fewer measurements are indicated. Capacity measurements are again proposed for one bridge (the Golden Gate Bridge) and one tunnel (the Caldecott tunnel, a multi-bore highway tunnel). These components are again similar to many other components in the system, but represent more valuable components, with higher failure costs than comparable components selected under the conditional entropy metric. This is the result of the inclusion of the economic factors of component management through the decision-making problem, as captured by the VoI metric.

Table 9-B: Sensor placements by the VoI metric, in order of greedy selection.

	Component Name	**Measure**
1	Caldecott Tunnel	Capacity
2	Caldecott Tunnel	Demand
3	Golden Gate Bridge	Capacity
4	San Francisco Oakland Bay Bridge (West Span)	Demand
5	Colden Gate Bridge	Demand

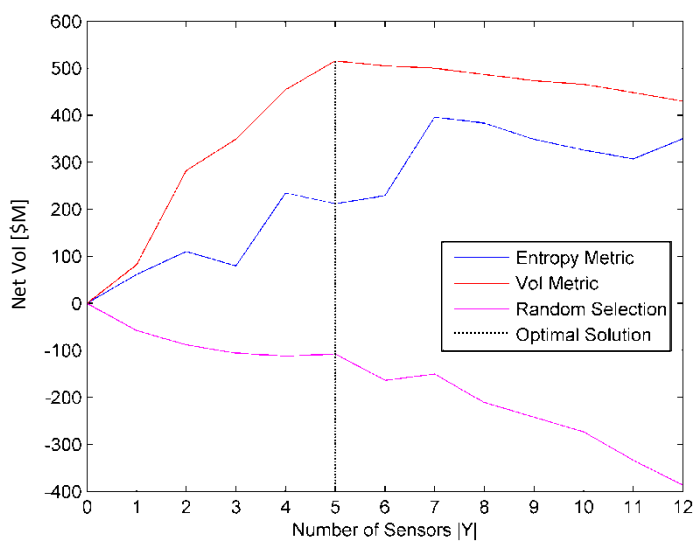

Figure 9-3: Net VoI versus number of sensors for sensing schemes optimized via the conditional entropy and VoI metrics, and for an example of random sensor placement.

As a quantitative comparison of these metrics, Figure 9-3 indicates the net VoI versus number of sensors for different measurement schemes, optimized greedily using the VoI or conditional entropy metrics, or using measurements selected at random. The scheme optimized by the VoI

metric clearly dominates schemes selected by other metrics. However, sensor placements selected using the conditional entropy metric outperform random selections, indicating that this metric might be used as a proxy for VoI in this case. Also note that, for random sensor placements, the net benefit of the sensing schemes is negative here, indicating that an uninformed approach to sensor placement for managing this system is inappropriate.

9.2.1. Value of Information Density

As a further investigation of the behavior of the VoI metric in this problem, the density of VoI across scenario parameters (i.e., earthquake magnitudes and epicenter locations) is evaluated for the proposed monitoring scheme of Figure 9-2. This density indicates which of these scenarios contribute most to the overall VoI of the scheme. A potential application of this investigation is to guide an alternative importance-sampling approach to the approximation of Equation 9-2.

Beginning with the formulation of Equation 9-2, and noting that the scenario parameter distribution p_W factorizes as the product of the distribution over magnitudes p_{W_m} and the distribution over epicenter locations p_{W_e}, the expectation over scenario parameters can be treated in two parts. The discounted cumulative VoI of Equation 9-4 can therefore be expressed as:

$$\mathrm{VoI}_{discounted\ cumulative}(Y) = \begin{cases} \int_{W_m} \Psi_Y(w_m)\ \mathrm{d}w_m \\ \int_{W_e} \Theta_Y(\mathbf{w}_e)\ \mathrm{d}\mathbf{w}_e \end{cases} \tag{9-5}$$

where the density of the discounted cumulative VoI with respect to earthquake magnitude, averaging over epicenter locations, is:

$$\Psi_Y(w_m) = \frac{\Lambda_{annual}}{\gamma_{annual}}\ \mathrm{p}_{W_m}(w_m)\ \mathbb{E}_{W_e}\mathrm{VoI}(Y|\mathbf{w}) \tag{9-6}$$

and the density of the discounted cumulative VoI with respect to epicenter location, averaging over magnitude, is:

$$\Theta_Y(\mathbf{w}_e) = \frac{\Lambda_{annual}}{\gamma_{annual}}\ \mathrm{p}_{W_e}(\mathbf{w}_e)\ \mathbb{E}_{W_m}\mathrm{VoI}(Y|\mathbf{w}) \tag{9-7}$$

Figure 9-4 plots the discounted cumulative VoI density with respect to earthquake magnitudes, both for the optimal sensor placement scheme depicted in Figure 9-2 as well as for a set consisting of all 54 possible measurements considered in this problem (representing the VoPI). Also depicted is the probability distribution over the earthquake magnitudes.

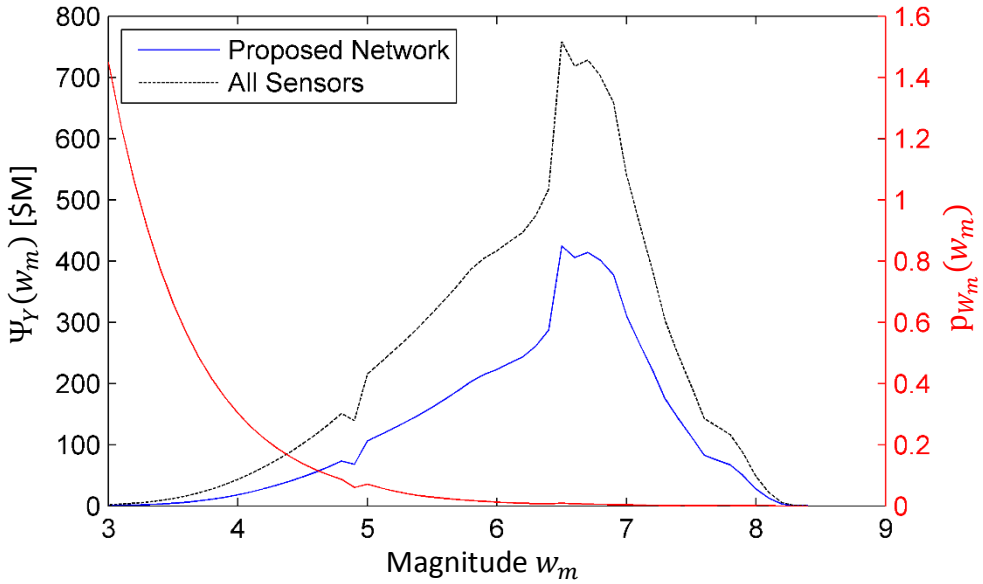

Figure 9-4: Density of the discounted cumulative VoI over earthquake magnitude for the proposed scheme of Figure 9-2 (blue), and for all possible sensors (black). The probability density function for earthquake magnitudes is also presented for reference.

Note that the density of the discounted cumulative VoI is low for both high-magnitude and low-magnitude events, with density peaking for moderate magnitudes. This can be understood both in terms of the destructive power of these earthquakes and their relative frequencies. Low-magnitude scenarios are common but are less damaging, so probabilities of component failure are low, making optimal management decision-making straightforward. VoI is therefore low in these scenarios, since additional information has little effect on the decision-making. For high-magnitude earthquakes, destructive power is high, and so the optimal management actions are again straightforward, i.e., to close down bridges and tunnels. VoI density is therefore low again, both because of this and because these events are relatively rare. In between these extremes, these effects balance out, with moderately common, moderately damaging events having less certain management decisions. Discounted cumulative VoI density is therefore highest between approximately magnitudes 6.5 and 7. Finally, note that taking the integrals under the curves of these figures shows that the optimal monitoring scheme achieves more than half of the VoPI while using less than a tenth of the potential measurements.

Examining the density of the discounted cumulative VoI with respect to epicenter locations, as shown in Figure 9-5 (plotting this density per kilometer along the region's fault lines), epicenter locations nearer to the center of the region contribute more heavily to the VoI of the optimal sensor network. There are several possible explanations of this result, including the clustering of both infrastructure components and sensors in this area. Also, the fault line with the highest density also has the highest probability for earthquake occurrence. All of these factors likely contribute to the high discounted cumulative VoI density along this fault.

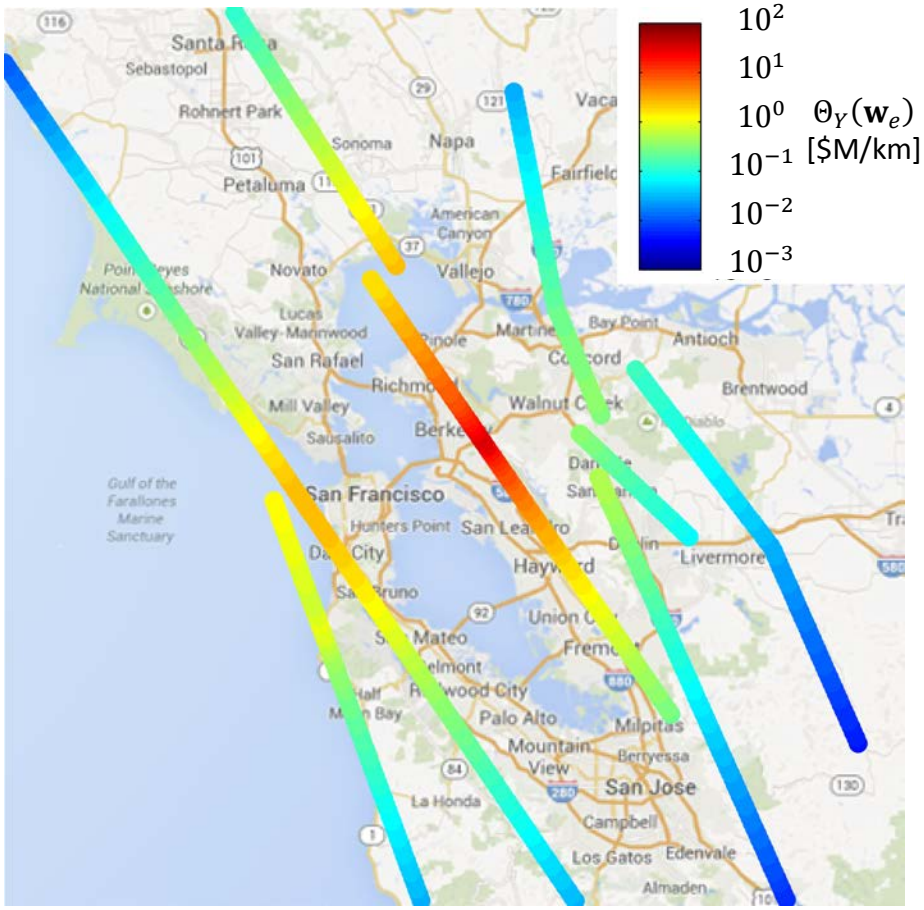

Figure 9-5: Density of the discounted cumulative VoI over earthquake epicenter locations for the proposed scheme of Figure 9-2. Note that the color scale is logarithmic. Background image from www.maps.google.com.

Finally, note that these results for discounted cumulative VoI density presented here refer only to the sensor network configuration presented in Figure 9-2. It is possible that different sensor network configurations will exhibit different VoI density distributions. The densities for the conditional entropy metric may appear different as well. Also note that these are densities for the VoI, not the expected loss due to these earthquakes. Scenarios which cause the highest loss to the system may not necessarily correspond to those with the highest VoI, since the latter measures the reduction in expected loss, not its magnitude.

Chapter 10

Urban Heat Risk Prediction and Response

This chapter discusses the application of the sensor placement methodology to the problem of temperature prediction in urban areas, to support decision-making on heat advisory issuance. Extreme high temperatures cause numerous direct (e.g., heat exhaustion, heatstroke) and indirect impacts (e.g., increased energy consumption) on human populations. These impacts are especially severe in urban areas, where a combination of higher temperatures due to the urban heat island effect and higher vulnerability due to urban demographics leads to a higher overall risk from extreme temperatures. To mitigate this risk, municipalities and their residents have several options, such as the issuance of targeted heat advisories and investments in green infrastructure and/or albedo reduction to reduce temperatures. To guide these mitigation activities, a comprehensive model of urban temperatures, supported by information for model improvement and updating, is necessary. For this reason, the sensor placement methodology developed here is applied to this problem. Further background information on heat risk in urban areas, urban microclimates, temperature prediction, and the options for mitigating extreme heat risk is provided in Section 10.1.

In this chapter, the random field being studied is the temperature (specifically, the surface temperature) in an urban area. As a first step, a spatio-temporal Gaussian probabilistic random field model is developed for this temperature. The characteristics and development of this model are presented in Section 10.2. Using this model, sensor placement is performed using the VoI metric, along with other metrics for comparative purposes. Various types of sensing activities are considered, from more traditional “point measurements” of temperature, e.g., thermometer readings, to pseudo-measurements of field components obtained from temperature forecasts and other sources. To define a decision-making problem in this context, a model for heat advisory issuance is developed that defines the components, states, actions, and losses considered in this problem. Details of observation activities and sensor placement metrics as applied in this problem are presented in Section 10.3.

In Section 10.4, the random field temperature model and various metrics are used to determine sensor placements for urban temperature monitoring. To illustrate this, a case study application is presented for Pittsburgh, Pennsylvania. Various observation types and prediction lead times are investigated, and the effects of these factors on the VoI provided by measurements are

investigated. Furthermore, the comparative performances of various metrics are investigated, illustrating how uncertainty reduction is modeled under these metrics.

As with Chapter 9, this chapter provides an illustrative development of a sensor placement optimization problem, handled under the methodology presented in this work. It begins with the development and calibration of a probabilistic model for the underlying random quantities affecting system performance. It then introduces a decision-making problem for the management of this system, i.e., a problem of heat-advisory issuance for avoiding the negative impacts of extreme heat. Finally, taking this model and decision-making problem, it illustrates the application of various sensor placement metrics, including VoI, to addressing temperature measurement for the representative urban area of Pittsburgh. The specific modeling assumptions adopted in this chapter do not necessarily reflect the priorities or objectives of relevant decision-makers and stakeholders, and therefore the results presented here should not be interpreted as a prescription for heat risk monitoring or mitigation in this city.

Results presented in this chapter were developed under the auspices of the Surface Heat Assessment for Developed Environments (SHADE) project. This work was funded in part by the Dowd Fellowship from the College of Engineering at Carnegie Mellon University. The author would like to thank Philip and Marsha Dowd for their financial support and encouragement. This project was also partially supported by the Metro21 Initiative at Carnegie Mellon University. The goal of this initiative is to research, develop, and deploy novel solutions to 21st-century problems facing metropolitan areas.

10.1. Background on Urban Heat Risk and Response

Recent extreme heat waves across North America, Europe, and Australia, and their adverse health impacts (Andrews, 1994; Kalkstein and Greene, 1997; Geerts and Linacre, 1999; Davis et al., 2003; Robine et al., 2008; Berko et al., 2014; Wells and Klima, 2016) have prompted many cities to focus their attention on adaptation, hazard mitigation, and emergency response plans for high temperatures (Menne and Matthies, 2009; Parson and Jameson, 2012; New York City Panel on Climate Change, 2013; Sustainable DC, 2013; Rhoades et al., 2014). During development and implementation of these plans, assessments of current and future temperature risks to the urban population can be helpful.

Average near-surface air temperatures (at 2m above ground level) in the Northeast U.S. are expected to increase between 2 and 5°C by 2100 (Melillo et al., 2014). Furthermore, in what is known as the urban heat island (UHI) effect, urban air temperatures are on average 1 to 3°C warmer than their rural surroundings during daytime (Oke, 1997; Sailor, 2002; US EPA, 2012), and up to 12°C warmer at night (Oke, 1997). The UHI effect interacts synergistically with heat waves to produce extreme heat stresses in urban areas that exceed the simple sum of the two effects (Li and Bou-Zeid, 2013). To understand and forecast the associated risks, physical

models of urban microclimates have been developed and embedded in larger scale weather and climate models.

One such model is the Princeton Urban Canopy Model (PUCM) (Wang et al., 2013), developed as an extension of the National Center for Atmospheric Research's WRF-ARW (Weather Research and Forecasting – Advanced Research WRF) model (Skamarock et al., 2005). The PUCM extension incorporates fine resolution representations of urban surface heterogeneity and hydrology, significantly improving the urban temperature forecasting over WRF's default urban canopy model (Li and Bou-Zeid, 2014); the combined model is denoted WRF-PUCM. However, such deterministic models are computationally intensive, making them difficult to use for short-term predictions of urban temperatures, where they must be re-run as updated global meteorological information is obtained and under a variety of initial conditions to reflect a range of uncertainty (Worley et al., 2011; Lin-Jiong et al., 2012; Sper de Almeida and Bauer, 2012). While deterministic models can express uncertainty implicitly through variation of the model inputs among multiple runs, probabilistic modeling explicitly represents this uncertainty; for this reason, a Gaussian process probabilistic model of temperatures is developed for this application. For further background on Gaussian processes, the reader is referred to Section 3.3, as well as to Rasmussen and Williams (2006). This spatio-temporal probabilistic temperature model, outlined in Section 10.2, is used to quantify the hazard posed by extreme heat in an urban area.

It is well established that increasing temperatures exacerbate both heat-related and air-quality-related illnesses and deaths (e.g., Clougherty and Kubzansky, 2009), so in addition to modeling the temperature hazard, it is important to understand how interactions between hazard, exposure, and vulnerability to high temperature create risk. These interactions are typically quantified as a product of terms relating to the hazard posed by the risk factor, and the vulnerability and exposure of the population (e.g., Crichton, 1999). This approach has been used in the literature to define extreme heat risk using various metrics for hazard, vulnerability, and exposure (e.g., Grower et al., 2011; Buscail et al., 2012; Aubrecht and Özceylan, 2013). In cases where exposure is combined with vulnerability, this combined metric is multiplied with the hazard metric to quantify risk.

Vulnerability indicates a population's sensitivity to extreme heat, and is usually quantified in terms of mortality (e.g., Doyon et al., 2008). The relative risk of mortality due to extreme heat begins to increase at approximately 27°C (e.g., Curriero et al., 2002). There is typically a positive correlation between temperature and non-accidental hospital admissions (e.g., Hondula and Barnett, 2014), indicating an increase in both mortality and morbidity with higher temperatures. Many factors impacting vulnerability have been identified in the literature, e.g., age, socio-economic status, and whether individuals live alone (e.g., Brunner, 2001; Hajat et al., 2006; Bosello et al., 2007; Stafoggia et al., 2008; Rosenthal et al., 2014). Additional literature has used factor analysis to identify which factors contribute most to the variability of the set of all factors within a vulnerability index (e.g., Reid et al., 2009; Harlan et al., 2012; Aubrecht and Özceylan, 2013; Bradford et al., 2015).

Exposure of populations to extreme heat is difficult to quantify, both due to the movement of people throughout the day and to factors such as access to air conditioning, which are sometimes considered as influencing exposure (Grower et al., 2011). Because of this, previous work has used population thresholds as a proxy for exposure, defining a binary exposure category (e.g., Buscail et al., 2012) or has analyzed exposure and vulnerability as a combined metric (e.g., Aubrecht and Özceylan, 2013).

Vulnerability of the population to extreme heat is not a function of the weather and, while exposure might be altered during extreme heat waves (e.g., the population prefers to remain indoors), it is also not strongly dependent on meteorological conditions. On the other hand, the hazard component of the risk is a direct function of urban temperatures and fluctuates significantly in space and over time. Therefore, risk forecasting to trigger heat advisories and mitigation plans requires accurate predictive hazard models or indicators. Extreme heat hazard has been assessed in a variety of ways, including using the average near-surface air temperature (e.g., Grower et al., 2011), the surface temperature at a specific time during a heat wave obtained from satellite data (e.g., Buscail et al., 2012; Harlan et al., 2012), or the total number of heat wave days in a given year (e.g., Aubrecht and Özceylan, 2013). In this chapter, fine resolution simulations of historical extreme heat events are used to develop a probabilistic surrogate Gaussian process model of spatio-temporal heat hazard in an urban area, as discussed above. This model captures the properties of the UHI effect and quantifies the variability in surface temperatures around the prior model prediction. Using such a model, various hazard statistics (e.g., average temperatures, 95th percentile temperatures, expected frequency of exceeding a certain threshold) might be extracted. However, the complete probabilistic model provides a full characterization of the heat hazard; extracting statistics from this model would mean that only a portion of the model information is being used to characterize the hazard.

The practical aim of the work presented in this chapter is to support the various strategies available to urban decision-makers to mitigate the risks posed by extreme heat events. These include short-term response strategies, such as the issuance of heat advisories warning people to remain in cooler areas (e.g., Grower et al., 2011) or the opening of public buildings as cooling centers where individuals may seek shelter, thereby reducing exposure to extreme heat (e.g., Kisner et al., 2012). Long-term mitigation strategies are also possible, including the use of cool or green roofs (Li et al., 2014) or other green infrastructures (US EPA, 2012) to reduce the intensity of the UHI effect, thereby reducing the extreme heat hazard of future events. Various risk mitigation and response strategies have been proposed for urban areas in the United States and abroad (Menne and Matthies, 2009; Parson and Jameson, 2012; New York City Panel on Climate Change, 2013; Sustainable DC, 2013; Rhoades et al., 2014). Actions responding to extreme heat risk have intrinsic costs (relating to societal disruption and increased energy consumption) and benefits (reduced mortality and morbidity due to heat exposure). These actions can also have varying effectiveness, e.g., lack of public awareness of what to do in response to heat advisories can limit their effectiveness (Wells and Klima, 2016). For these reasons, the VoI

metric is applied in this chapter to address the problem of temperature monitoring to support heat risk mitigation decision-making in urban areas.

10.2. Probabilistic Temperature Modeling

This section describes the definition, calibration, and validation of a spatio-temporal Gaussian random field model for urban temperatures, as well as what measurements for model updating are considered, and how these measures are processed. For model validation, two case-study cities are investigated: New York City, New York, and Pittsburgh, Pennsylvania.

10.2.1. Definition

The random temperature field is first decomposed into three elements. Each element constitutes a separate random field, and each is described using a Gaussian process random field model. Section 10.2.1.1 provides a description (using standard meteorological notation) of the temperature decomposition. Sections 10.2.1.2 through 10.2.1.4 discuss the specific models developed for the global average, cyclic pattern, and residual elements of the temperature respectively. In Section 10.2.1.5, these models are combined into the overall Gaussian process temperature model.

10.2.1.1. Temperature decomposition

Given a location $\mathbf{x}$ and a time t, let the spatio-temporal field of surface temperature over a specified domain (in degrees Celsius) be denoted by $\mathrm{T}(\mathbf{x}, t)$. This temperature field is decomposed as:

$$\mathrm{T}(\mathbf{x}, t) = \mathrm{T}_0(t) + \mathrm{T}'(\mathbf{x}, t) \tag{10-1}$$

where $\mathrm{T}_0(t)$ is the mean surface temperature over the domain (i.e., over all $\mathbf{x} \in X$, for each time $t \in \tau$) and $\mathrm{T}'(\mathbf{x}, t)$ is the perturbation from this value (note that the expected value of this perturbation over $\mathbf{x}$ is zero at each time t). Given this formulation, $\mathrm{T}_0(t)$ can be interpreted as the average temperature over a modeled region of interest, e.g., the average city-wide temperature as a function of time.

The perturbation is further decomposed as:

$$\mathrm{T}'(\mathbf{x}, t) = \mathrm{T}_1(\mathbf{x}, t) + \mathrm{T}''(\mathbf{x}, t) \tag{10-2}$$

where $\mathrm{T}_1(\mathbf{x}, t)$ is the spatially and temporally varying temperature pattern at location $\mathbf{x}$ and at time t. This pattern is defined to be cyclic, recurring at a fixed time interval Δt, i.e., $\mathrm{T}_1(\mathbf{x}, t + q\Delta t) = \mathrm{T}_1(\mathbf{x}, t)$ for any integer q. A cycle length of $\Delta t = 24$ hours is used, capturing a daily cyclic component of the spatial distribution of the temperature perturbation. In the context of urban temperatures, these distributions capture the UHI effect as it varies across space and

throughout the day. Depending on the application, additional terms may be added to include other cyclic temperature patterns, e.g., seasonal patterns. Finally, $\mathrm{T}''(\mathbf{x},t)$ is the temperature residual, which captures the variability in the temperature that is not described by either $\mathrm{T}_0(t)$ or $\mathrm{T}_1(\mathbf{x},t)$. The origins of $\mathrm{T}''(\mathbf{x},t)$ are stochastic and can be traced to various physical processes such as anthropogenic heat emissions (from cars or heating/cooling systems) or the influence of wind patterns and clouds on surface temperatures.

Combining Equations 10-1 and 10-2 above yields the full temperature decomposition:

$$\mathrm{T}(\mathbf{x},t) = \mathrm{T}_0(t) + \mathrm{T}_1(\mathbf{x},t) + \mathrm{T}''(\mathbf{x},t) \tag{10-3}$$

10.2.1.2. Global mean temperature

The global mean of the temperature $\mathrm{T}_0(t)$ for the region at time t is modeled as:

$$\mathrm{T}_0(t) \sim \mathcal{GP}\left(\mu_{T_0}, \mathrm{k}_{T_0}(t_i,t_j)\right) \tag{10-4}$$

where μ_{T_0} denotes the expected mean temperature in the region over time and $\mathrm{k}_{T_0}(t_i,t_j)$ models the temporal covariance in these temperatures. The average μ_{T_0} is estimated from summertime temperature records simulated by the WRF-PUCM model as training data for a region (as discussed in Section 10.2.3.2). This value will in fact vary seasonally and over longer time horizons due to climate change. However, as this model is used for predictions over periods of several hours, variation in this term is assumed to be negligible.

A covariance function with a combined exponential and periodic correlation structure is used to model the daily cyclic variations in global temperatures. This function is:

$$\mathrm{k}_{T_0}(t_i,t_j) = \sigma_{T_0}^2\left(\alpha_{T_0}\exp\left(-\frac{|t_i-t_j|}{\lambda_{T_0}}\right) + \left(1-\alpha_{T_0}\right)\cos\left(2\pi\frac{|t_i-t_j|}{\Delta t}\right)\right) \tag{10-5}$$

This function models the decay in correlation between the average temperatures at two times (t_i and t_j) as the absolute difference between these times, $|t_i - t_j|$, increases. It also models the periodic correlation, with a period of $\Delta t = 24$ hr, capturing the day/night temperature cycle. The parameter λ_{T_0} is the correlation timescale, σ_{T_0} is the marginal standard deviation of the average temperature, and α_{T_0} is a weighting coefficient between the two parts of the covariance function (i.e., if $\alpha_{T_0} = 1$, the correlation function is purely exponential, while if $\alpha_{T_0} = 0$, it is purely periodic). The evaluation of these parameters from high-resolution dynamic simulations is described in Section 10.2.3.2.

10.2.1.3. Cyclic temperature pattern

The cyclic temperature pattern captures the systematic differences in temperature between different parts of a region and the global average temperature at different times of the day. This pattern is modeled as a spatio-temporal Gaussian process:

$$\mathrm{T}_1(\mathbf{x}, t) \sim \mathcal{GP}\left(0, \mathrm{k}_{T_1}\left(\mathbf{x}_i, t_i, \mathbf{x}_j, t_j\right)\right) \tag{10-6}$$

The expected value of this pattern in space and time is assumed to be zero. There will in fact be areas that are systematically above or below the average temperature, e.g., due to the UHI effect of the spatial variation of the properties of the urban fabric. However, assuming a prior mean of zero allows these patterns to be learned from training data, and thus captured in the posterior mean field of the cyclic temperature pattern.

The covariance function is further assumed to be separable into a spatial and a temporal covariance:

$$\mathrm{k}_{T_1}\left(\mathbf{x}_i, t_i, \mathbf{x}_j, t_j\right) = \sigma_{T_1}^2 \mathrm{k}_{T_1,X}\left(\mathbf{x}_i, \mathbf{x}_j\right) \mathrm{k}_{T_1,\tau}\left(t_i, t_j\right) \tag{10-7}$$

The parameter $\sigma_{T_1}^2$ describes the variance of the cyclic pattern. The (normalized) spatial covariance is:

$$\mathrm{k}_{T_1,X}\left(\mathbf{x}_i, \mathbf{x}_j\right) = \alpha_{T_1,X} \exp\left(-\frac{|\mathbf{x}_i - \mathbf{x}_j|}{\lambda_{T_1,X,1}}\right) + \left(1 - \alpha_{T_1,X}\right) \exp\left(-\frac{|\mathbf{x}_i - \mathbf{x}_j|}{\lambda_{T_1,X,2}}\right) \tag{10-8}$$

This represents the weighted average of two exponential correlation functions with different length parameters $\lambda_{T_1,X,1}$ and $\lambda_{T_1,X,2}$, with weighting coefficient $\alpha_{T_1,X}$. This mixture of two functions was found to perform the best in capturing the empirical correlation structure, as described in Section 10.2.3.3, with the fewest parameters.

The (normalized) temporal covariance is:

$$\mathrm{k}_{T_1,\tau}\left(t_i, t_j\right) = \left(1 - \alpha_{T_1,\tau}\right) \exp\left(-\frac{|t_i - t_j|_{\mathrm{cyl}}}{\lambda_{T_1,\tau}}\right) + \alpha_{T_1,\tau} \tag{10-9}$$

where $\left|t_i - t_j\right|_{\mathrm{cyl}}$ denotes the cyclic time duration between t_i and t_j:

$$\left|t_i - t_j\right|_{\mathrm{cyl}} = \min\left(\left|t_i - t_j\right| \bmod \Delta t, \Delta t - \left|t_i - t_j\right| \bmod \Delta t\right) \tag{10-10}$$

where $\Delta t = 24$ hr to describe a daily cycle. For example, the cyclic time between 8PM and 4AM is 8 hr. The temporal covariance therefore reflects an exponential correlation applied to a transformation of the temporal domain to give the covariance a cyclic property. The use of this cyclic exponential model, rather than a sinusoidal model as in Equation 10-5, better matches the

empirical correlation structure of the cyclic temperature patterns, as discussed in Section 10.2.3.3.

Note that the model for temporal covariance of these patterns does not include a continuous decay, as does the model for T_0. This encodes the assumption that these cyclic patterns are stationary over time for each part of the day. This assumption allows historical simulation data to be used for assessing the spatial distribution of temperatures for future model predictions. Changes in land cover, including the deployment of cool or green roofs on buildings, alter the UHI pattern of a region, but they can be modeled by fine resolution simulations (Li et al., 2014). Therefore, periodic re-calibration of these the cyclic patterns should be conducted to reflect changes in land usage and capture the impacts of these changes on the UHI. Alternatively, by training these patterns indirectly using a dense network of temperature measurements over the region of interest along with the full probabilistic temperature model of Section 10.2.1.5, the patterns can be periodically re-calibrated as new data are obtained.

10.2.1.4. Residual temperature

The residual temperature field captures all temperature phenomena not included in the global average and cyclic pattern models above. This includes the effects of factors such as solar irradiation, soil moisture, wind, and cloud cover on surface temperature. The influence of these factors is deterministically modeled by WRF-PUCM, at a relatively high computational cost. In the probabilistic model, however, the effects of these factors are treated as sources of randomness in the temperature field about its expected value at a particular space-time coordinate, and this randomness is captured in the residual temperature term $\mathrm{T}''(\mathbf{x}, t)$. This term is modeled as a Gaussian process:

$$\mathrm{T}''(\mathbf{x}, t) \sim \mathcal{GP}\left(0, \mathrm{k}_{T''}\left(\mathbf{x}_i, t_i, \mathbf{x}_j, t_j\right)\right) \tag{10-11}$$

The mean over space and time of this residual is zero by definition, since any recurring pattern in the residual $\mathrm{T}'(\mathbf{x}, t)$ has been included in the $\mathrm{T}_1(\mathbf{x}, t)$ term. The covariance is again assumed to be separable between space and time:

$$\mathrm{k}_{T''}\left(\mathbf{x}_i, t_i, \mathbf{x}_j, t_j\right) = \sigma_{T''}(\mathbf{x}_i)\sigma_{T''}\left(\mathbf{x}_j\right)\mathrm{k}_{T'',\mathrm{X}}\left(\mathbf{x}_i, \mathbf{x}_j\right)\mathrm{k}_{T'',\tau}\left(t_i, t_j\right) \tag{10-12}$$

The variance of the residual is modeled as a function of its spatial coordinate, i.e., $\sigma_{T''}(\mathbf{x})$ denotes the standard deviation of the residual field at location $\mathbf{x}$. The spatial covariance function is the (normalized) exponential covariance:

$$\mathrm{k}_{T'',\mathrm{X}}\left(\mathbf{x}_i, \mathbf{x}_j\right) = \exp\left(-\frac{|\mathbf{x}_i - \mathbf{x}_j|}{\lambda_{T'',X}}\right) \tag{10-13}$$

The temporal covariance function is the weighted average of two exponential functions:

$$\mathrm{k}_{T'',\tau}(t_i,t_j) = \alpha_{T'',\tau}\exp\left(-\frac{|t_i-t_j|}{\lambda_{T'',\tau,1}}\right) + \left(1-\alpha_{T'',\tau}\right)\exp\left(-\frac{|t_i-t_j|}{\lambda_{T'',\tau,2}}\right) \tag{10-14}$$

This mixture model best captures the empirical correlation of the residual temperature field, as discussed in Section 10.2.3.4.

10.2.1.5. Gaussian process temperature model

Assuming that the models of the temperature elements described above are mutually independent, they can be combined to define a prior Gaussian process model of the temperature field as:

$$\mathrm{T}(\mathbf{x},t) \sim \mathcal{GP}\left(\mu_{T_0}, \mathrm{k}_{T_0}(t_i,t_j) + \mathrm{k}_{T_1}(\mathbf{x}_i,t_i,\mathbf{x}_j,t_j) + \mathrm{k}_{T''}(\mathbf{x}_i,t_i,\mathbf{x}_j,t_j)\right) \tag{10-15}$$

This model allows the definition, for a finite set of locations X and finite time duration τ, of a prior multivariate Gaussian distribution for $\mathbf{T}$, the vector of temperatures at of all pairs $\{\mathbf{x},t\}$ of space and time coordinates in sets X and τ. This distribution is:

$$\mathbf{T} \sim \mathcal{N}(\boldsymbol{\mu}_T, \boldsymbol{\Sigma}_T) \tag{10-16}$$

where $\boldsymbol{\mu}_T$ is the mean vector and $\boldsymbol{\Sigma}_T$ is the covariance matrix.

10.2.2. Measurements

Given any observation $\mathbf{y}$ of this temperature field (or any of its parts, such as the global mean temperature), the prior model of Section 10.2.1.5 can be updated to a posterior model. Three types of observations are considered for model updating: measurements of the average temperature in the modeling region (obtained as pseudo-measurements from temperature forecasts), assessments of the cyclic temperature patterns (obtained by processing simulation data on the historical urban microclimate), and local temperature measurements (based on field observations). Combining data from these various sources into a single observation vector $\mathbf{y}$, the prior temperature field model can be updated to a posterior model, and the resulting temperature prediction accuracy can be improved.

10.2.2.1. Global mean temperature measurements

The Gaussian process model for the global average temperature discussed in Section 10.2.1.2 can be used for prediction of future average temperatures based on past data or on future forecasts from other sources. Such data are considered to be observations of the global average of temperatures in a region, i.e.:

$$\mathrm{y}_{T_0}(t) = \frac{1}{n_X}\sum_{\mathbf{x}\in \mathrm{X}} \mathrm{T}(\mathbf{x},t) + \epsilon_{T_0}(t) \tag{10-17}$$

Errors in these measurements are denoted $\epsilon_{T_0}(t)$, and are in general modeled as Gaussian random variables. Here it is assumed that this error is negligible, i.e., $\epsilon_{T_0}(t) = 0 \; \forall \; t \in \tau$, since the average temperatures used for this model are computed from simulated temperature data (as discussed in Section 10.2.3.1). However, it should be noted that the use of different simulation domain definitions and/or spatial resolutions between the mean temperature observations and the Gaussian process model can lead to a systematic bias of these measurements of the global average. Such bias can be accounted for in the Gaussian process framework, i.e., by defining the joint distribution for $\epsilon_{T_0}(t)$ with correlated noise between measurements at different times.

Observations of the mean temperature might be physical measurements, obtained by averaging many local observations over the region, or pseudo-observations, obtained from coarse resolution weather or climate models where the average temperature over the entire region is the output of a single cell in the model. In this way, different types of data can be integrated to update model predictions. Note that a vector of global mean temperature measurements at a finite set of times $\mathbf{y}_{T_0}$ can be expressed in the form of Equation 3-5, i.e., that these observations are a linear combination of the discretized temperature field $\mathbf{T}$ (through appropriate definition of $\mathbf{\Omega}_Y$) and vector of observation errors $\boldsymbol{\epsilon}_{T_0}$.

For the application purposes described in this chapter, regional forecasts of the average temperature obtained from coarse resolution weather simulations for the region in question are processed as pseudo-observations of the global mean, as in Equation 10-17. For example, forecasts by the National Weather Service (NWS) made using coarse resolution WRF models can be incorporated as pseudo-observations. The result is an updated posterior model for the global average temperature with much less variance than the prior model, assuming low error in these coarse resolution model predictions. Models of the cyclic temperature patterns and residuals are then used to perform statistical downscaling from these coarse resolution predictions to improve fine resolution forecasting.

10.2.2.2. Cyclic temperature pattern measurements

In determining cyclic temperature patterns, available fine resolution data from historical events can provide pseudo-observations for updating the prior model. Such observations are of the form:

$$\mathrm{y}_{T_1}(\mathbf{x}, t_{\mathrm{cyl}}) = \frac{1}{n_{\mathrm{cyl}}} \sum_{t \in \tau_{\mathrm{cyl}}} \left[\mathrm{T}(\mathbf{x}, t) - \mathrm{y}_{T_0}(t)\right] + \epsilon_{T_1}(\mathbf{x}, t_{\mathrm{cyl}}) \tag{10-18}$$

where $\tau_{\mathrm{cyl}} = \left\{t_1, \dots, t_{n_{\mathrm{cyl}}}\right\}$ is the set of n_{cyl} time steps in the data set at the same time of day, such that $t \bmod \Delta t \; = \; t_{\mathrm{cyl}} \; \forall \; t \in \tau_{\mathrm{cyl}}$. That is, these pseudo-observations represent the average difference between the temperature field and the global average temperature (measured as in Equation 10-17) at a certain time during the diurnal cycle. Again, the error in these measurements, $\epsilon_{T_1}(\mathbf{x}, t_{cyl})$, is assumed to be negligible, i.e., $\epsilon_{T_1}(\mathbf{x}, t_{\mathrm{cyl}}) = 0 \; \forall \; \mathbf{x} \in X, t_{\mathrm{cyl}} \in \tau_{\mathrm{cyl}}$.

Errors in these measurements resulting from insufficient data or lack of recent simulation results can be accounted for with appropriate definition of Gaussian error $\epsilon_{T_1}(\mathbf{x}, t_{\text{cycle}})$ in these pattern measurements. Note again these measurements can be expressed in the form of Equation 3-5 with an appropriate definition for $\mathbf{\Omega}_Y$. By conditioning this cyclic pattern on fine resolution temperature simulation data, information about the land characteristics of the region and their effect on the spatial distribution of temperature is captured by the posterior probabilistic model.

10.2.2.3. Local temperature measurements

Direct observations of the local temperature at location $\mathbf{x}$ and time t are modeled as:

$$\mathrm{y}_T(\mathbf{x}, t) = \mathrm{T}(\mathbf{x}, t) + \epsilon_T(\mathbf{x}, t) \quad \epsilon_T(\mathbf{x}, t) \sim \mathcal{GP}\left(\mu_\epsilon(\mathbf{x}, t), \mathrm{k}_\epsilon\left(\mathbf{x}_i, t_i, \mathbf{x}_j, t_j\right)\right) \tag{10-19}$$

That is, the measurement is the true local temperature at the time of the measurement, corrupted by a measurement error. This error can itself be defined with a Gaussian process model. In this chapter, a white-noise model is used, with zero-mean independent measurement error, i.e., $\mu_\epsilon(\mathbf{x}, t) = 0$ and $\mathrm{k}_\epsilon\left(\mathbf{x}_i, t_i, \mathbf{x}_j, t_j\right) = \delta_{ij}\sigma_\epsilon^2$ for any measurement, where δ_{ij} is the Kronecker delta. Note that these measurements are also expressible in the form of Equation 3-5.

10.2.3. Calibration

The computational complexity of fine resolution physical models of urban microclimates generally precludes their use for fast short-term temperature forecasting. Instead, the outputs of these models are used here to train a spatio-temporal probabilistic model, which allows the use of relatively fast probabilistic inference methods for temperature prediction, using the measurements considered in Section 10.2.2. This section describes the generation of training data from fine resolution models for New York City, NY and Pittsburgh, PA, USA and the use of these data in calibrating the parameters of the Gaussian process models discussed in Section 10.2.1.

In this section, a least-squares curve-fitting approach is used to match the empirical covariance derived from statistical analysis of the training data to the parametric forms for the covariance functions. Typically, calibration of model parameters is performed using maximum likelihood estimation, where the likelihood is computed over the training data set using the Gaussian process model with the proposed parameters (Rasmussen and Williams, 2006). However, in the application discussed here, calibration of the parameters via this method is computationally prohibitive due to the number of parameters involved and for the size of the training data sets (including more than 4 million samples each). Furthermore, rather than jointly calibrating all parameters, the parameters of each part of the full Gaussian process model are trained separately, and these trained models are combined to create the full model, as discussed in Section 10.2.1.5. Again, this is done to reduce the computational complexity of the calibration process. In

principle, any appropriate method for Gaussian process model calibration can be applied to this problem, and these may yield slightly different calibration results than those presented here.

10.2.3.1. Training data simulation using WRF-PUCM

To calibrate a probabilistic random field for temperatures, it is first necessary to obtain representative spatio-temporal temperature data for the region of interest. Training data simulated from the WRF-ARW model with the PUCM extension are used for this purpose. The simulations use the North American Land-use Category Dataset (NLCD) 2006 to determine the land use type and the surface properties. The boundary conditions that drive the simulations in the largest computational domain, as well as the initial conditions, are obtained from the North American Regional Reanalysis data at 6 hr intervals. This largest computational domain has a horizontal resolution of 9 km, and two additional domains with resolutions of 3 and 1 km are embedded therein to refine the grid over the target city. This grid refinement uses only one-way coupling: the larger domain provides the boundary conditions to drive the subsequent finer one, but the smaller domain does not provide feedback to the coarser one. All domains use 60 vertical grid levels, which are centered on either New York City, NY, or Pittsburgh, PA, for each of the regions in question.

The WRF model defaults to the following physical parameterization schemes: (i) the rapid radiative transfer model scheme for longwave radiation (Mlawer et al., 1997); (ii) the Dudhia scheme (Dudhia, 1989) for shortwave radiation; (iii) the 2D Smagorinsky scheme for horizontal diffusion; (iv) the mosaic Noah land surface model for non-urban surfaces; and (v) the Mellor−Yamada−Janjic planetary boundary layer scheme (Mellor and Yamada, 1974) along with the modified Zilitinkevich relationship for thermal roughness length parameterization (Chen and Zhang, 2009). Cumulus parameterization schemes are not used in the simulations since they are not needed with the fine resolutions adopted here, as the convective scales are dynamically resolved.

The default single layer urban canopy model (SL-UCM) inside WRF is replaced by the Princeton Urban Canopy Model (Wang et al., 2013). The PUCM, akin to SL-UCM, is based on the energy exchange framework developed by Kusaka et al. (2001). However, each facet (e.g., wall, roof, or road) in the PUCM can be further subdivided into multiple sub-facets. For example, the user can specify fractions of white/black/green roofs, brick/concrete walls, and concrete/asphalt/vegetated grounds. The multiple sub-facets can be modeled with distinct physical and thermal properties. Another important improvement in PUCM is the adoption of more realistic representations for hydrological processes, including in-canyon vegetated soils and water storage capacity for impervious materials. The PUCM coupled to WRF has been previously tested and produced results that better capture the UHI effect compared to the default UCM (Li and Bou-Zeid, 2014). The accurate fine-scale representation provided by this model allows the correlation scales of the trained Gaussian process model to better reflect the true scale of variation of urban temperatures.

Apart from replacing the default UCM, the WRF simulations also use a mosaic-based approach

(Li and Bou-Zeid, 2014) to compute the surface fluxes as opposed to the dominant category method commonly used in the default WRF-ARW model. In the mosaic-based approach, instead of fixing a single land use category at each grid point as the dominant category in that grid cell, the fluxes from the most common n_{types} number (n_{types} can be varied from 1 to 15 and added as an input in the namelist file of WRF) of land use types in that grid cell are solved and then fractionally summed to compute the fluxes into the atmosphere. The mosaic approach is pertinent for urban modeling as the land use type in most urban areas, even at 1 km grid spacing, is highly varied and there might not necessarily be a dominant land type category. In NLCD 2006, urban areas fall under three categories: low intensity (50% built cover), medium intensity (90% built cover) and high intensity (95% built cover). For example, the land use type for a grid cell inside the borough of Manhattan is as follows: 39% high intensity urban, 22% medium intensity urban, 20% low intensity urban, 10% green cover and 9% water. In the default approach, WRF would assume most tiles are high intensity urban. In the mosaic-based approach, WRF will solve for all five land use categories and will fractionally add the computed fluxes. While the default approach is appropriate for regional scale modeling (36 km - 10 km), the mosaic approach is more suited at much finer resolutions focusing on cities where the variability length scale of the surface is sharply reduced. The mosaic approach also improves the representation of fluxes from various land use categories without the need to further increase the model resolution, thereby cutting the computational cost. These modifications have been thoroughly tested and validated by Ramamurthy et al. (2015). Again, this fine resolution is important to allow the calibrated Gaussian process temperature model to accurately reflect the true length-scale of urban temperature variability.

Using this deterministic temperature model, simulations were conducted of summertime temperatures in both New York City and Pittsburgh to obtain data to calibrate and to validate the Gaussian process temperature model. A pair of simulations was conducted for each city, one for a shorter duration to provide data used to calibrate the Gaussian process model parameters, and the other to provide data for model validation, as discussed in Section 10.2.4. Table 10-A summarizes the simulation areas, spatial grid scales, start and end dates, durations, and temporal resolutions of the simulations. Note that where the calibration and validation periods overlap, the same simulation results are obtained.

Table 10-A: WRF-PUCM simulation settings.

Setting		Value (New York)	Value (Pittsburgh)
simulation area	[km]	159×159	198×201
spatial grid spacing	[km]	1	1
calibration period start date		July 15, 2006	June 15, 2012
calibration period end date		July 21, 2006	June 21, 2012
calibration period duration	[days]	7	7
calibration period temporal grid spacing	[hr]	1	0.5
validation period start date		July 15, 2006	June 15, 2012
validation period end date		August 9, 2006	July 13, 2012
validation period duration	[days]	25	29
validation period temporal grid spacing	[hr]	0.5	0.5

Figure 10-1 shows the simulated regions around New York City and Pittsburgh for reference. Most major distinctions between rural (green), urban (grey), and water (blue) areas are visible in these satellite images.

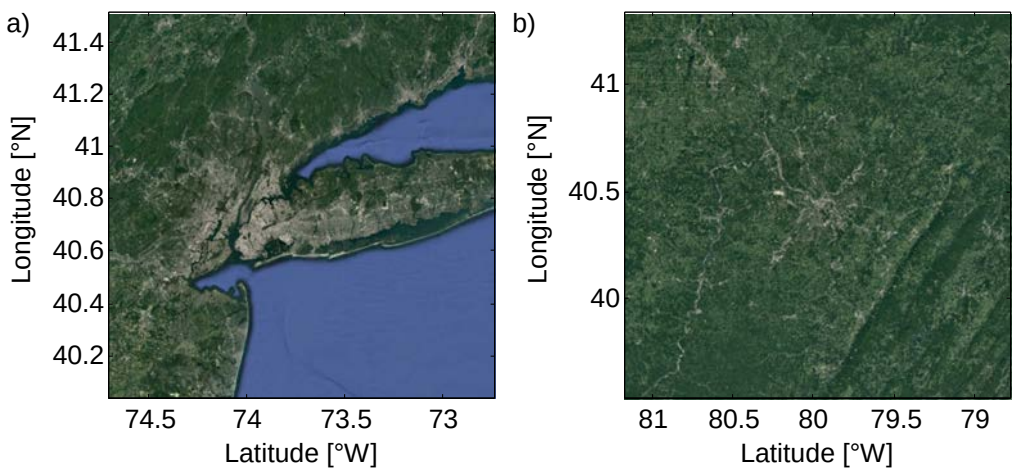

Figure 10-1: Satellite imagery of the simulation regions for New York City (a) and Pittsburgh (b). Satellite images from www.maps.google.com.

10.2.3.2. Global mean temperature

Figure 10-2 presents a comparison of the covariance model for the global mean temperature $T_0(t)$ from Equation 10-5, calibrated from data simulated for New York City, and the observed temporal covariance pattern in the calibration data. The strong cyclic pattern in the correlation structure reflects the daily temperature cycle of the data.

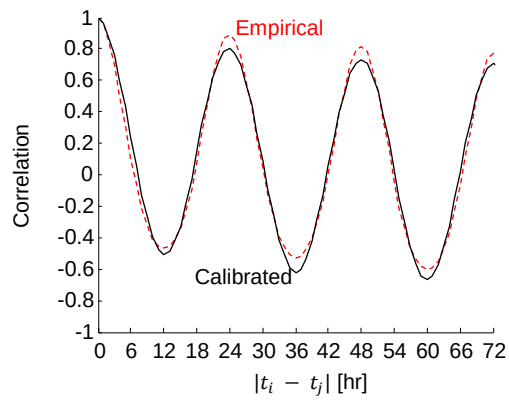

Figure 10-2: Comparison between the empirically observed temporal correlation of the global mean temperature, as computed from simulated temperature data for New York City (dotted line), and the calibrated covariance function of Equation 10-5 (solid line).

For reference, Table 10-B lists the parameters of this calibrated model for New York City and for Pittsburgh. Note that the average temperatures μ_{T_0} are calibrated for these cities based on simulated training data only, i.e., that these represent the average temperatures for these cities during a week of July or June, respectively, rather than year-round average temperatures.

Table 10-B: Calibrated parameters of the Gaussian process model for the global mean temperature.

Parameter		**Value (New York)**	**Value (Pittsburgh)**
μ_{T_0}	[°C]	31.5	24.5
σ_{T_0}	[°C]	5.97	6.66
α_{T_0}	[-]	0.31	0.34
λ_{T_0}	[hr]	23.7	23.3

Figure 10-3 presents an example of average temperature prediction making use of this model and global average temperature data for the simulated temperature record in New York City. Note that confidence in the forecast decreases as the prediction lead time increases. However, the daily periodicity of the temperature patterns is preserved, due to the cyclic correlation structure.

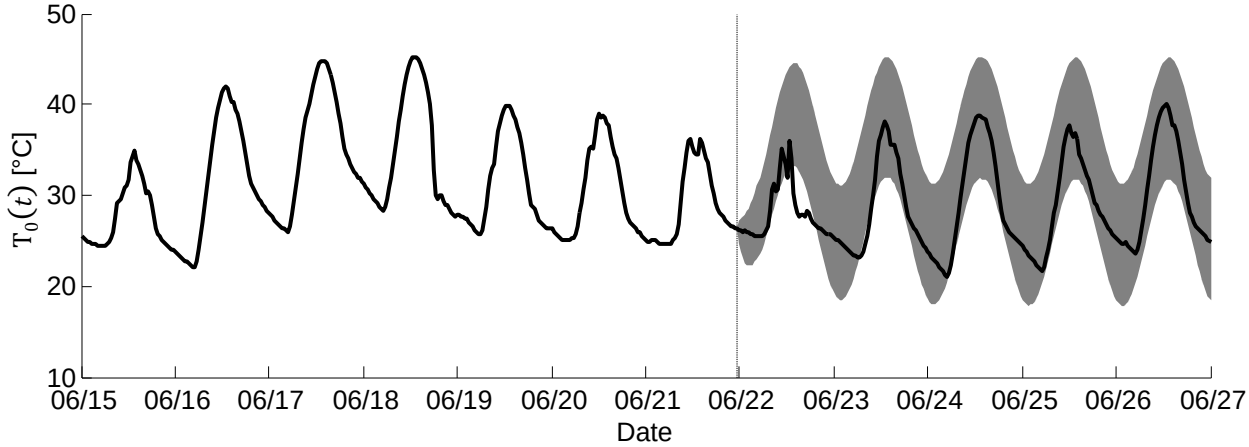

Figure 10-3: Temperature predictions using a Gaussian process model for the global mean temperature. The average temperatures before 6/22 are used as training data, and the resulting 95% confidence bounds on the predictions of future temperatures are shown as the grey area, along with the actual temperature record (black line).

10.2.3.3. Cyclic temperature pattern

Figure 10-4 compares calibrated covariance functions for the cyclic temperature pattern in space from Equation 10-8 and in time from Equation 10-9 to the empirical covariance observed in simulated temperature data from New York City, NY. Note that the empirical covariance is obtained after removing the global mean temperature $T_0(t)$ from the data.

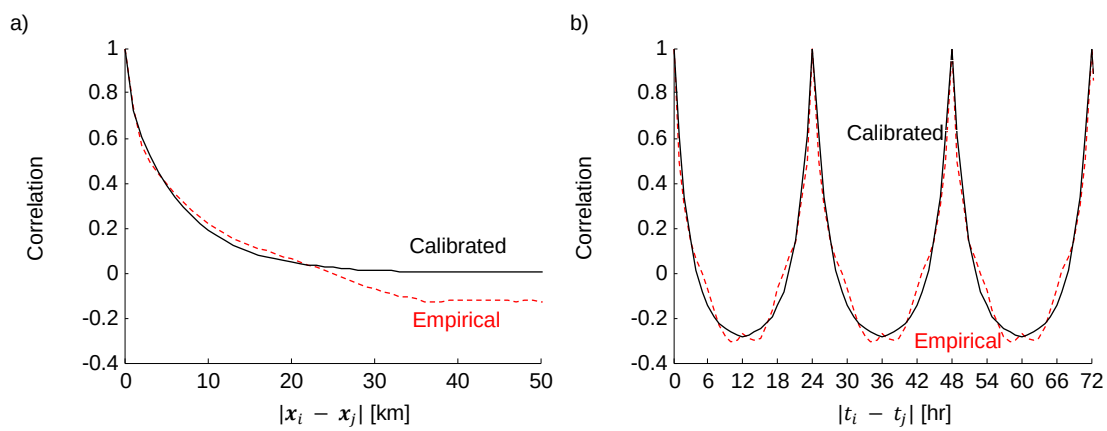

Figure 10-4: Comparison between the empirically observed spatial (a) and temporal (b) correlations of the cyclic temperature pattern, as computed from simulated temperatures for New York City (dotted line), and the calibrated covariance functions of Equations 10-8 and 10-9 (solid line).

For reference, Table 10-C lists the parameters of this calibrated model for New York City and for Pittsburgh. Note that, although the length-scale parameters $\lambda_{T_1,X,1}$ and $\lambda_{T_1,X,2}$ differ by an order of magnitude for New York, they are nearly the same for Pittsburgh. This indicates that a correlation model with a single decay term (rather than the two decay terms of Equation 10-8) would probably be adequate to model the spatial correlation structure for the cyclic temperature pattern in Pittsburgh. However, the model of Equation 10-8 is still used here for consistency.

Table 10-C: Calibrated parameters of the covariance function for the cyclic temperature pattern.

Parameter		Value (New York)	Value (Pittsburgh)
σ_{T_1}	[°C]	0.44	0.58
$\alpha_{T_1,X}$	[-]	0.8	0.7
$\lambda_{T_1,X,1}$	[km]	7.0	7.9
$\lambda_{T_1,X,2}$	[km]	0.5	7.8
$\alpha_{T_1,\tau}$	[-]	-0.3	-1.5
$\lambda_{T_1,\tau}$	[hr]	2.8	10.3

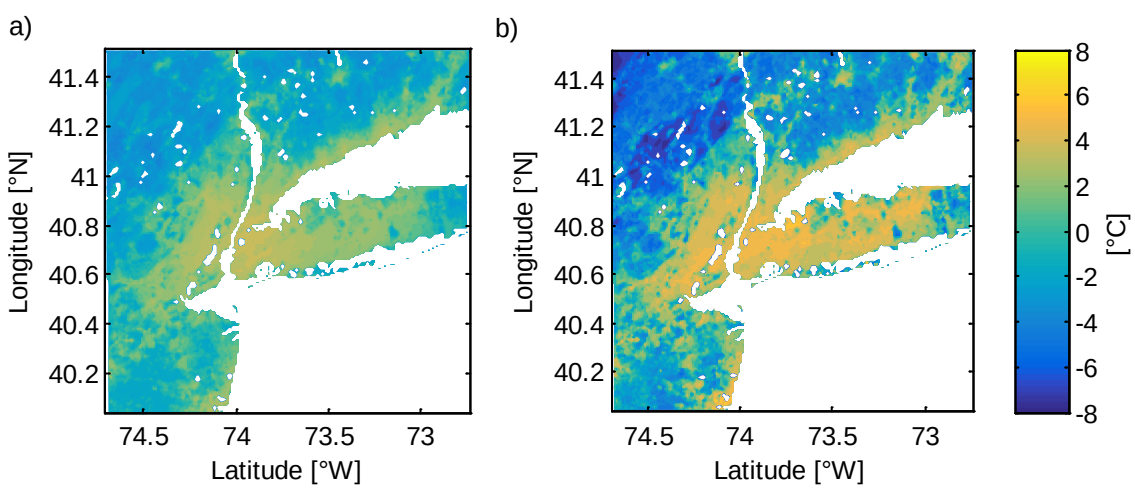

Figure 10-5: Posterior mean of the cyclic temperature patterns for New York City at local midnight (a) and noon (b), conditioned on measurements of these patterns, as in Equation 10-18, obtained from the training data set.

Figure 10-5 shows the trained cyclic patterns in the statistical model, i.e., the posterior mean field of the cyclic temperature pattern conditioned on training data, for New York City at local midnight and noon. Note that the UHI effect pattern for surface temperatures is more pronounced during the day than at night, which is as expected (Oke, 1997). Also note that these patterns are not considered to be part of the calibrated model; rather, these represent the mean field of the posterior Gaussian process model, conditioned on measurements y_{T_1} of the WRF-simulated data

during the calibration period. As such, these patterns are only considered to be "known" to the model if measurements y_{T_1} are available, as discussed in Section 10.2.4. Otherwise, these patterns must be learned from other data sources, i.e., from the local measures y_T of Equation 10-19.

10.2.3.4. Residual temperature

Figure 10-6 compares the empirical correlations of the residual temperature field for New York City and the fitted correlation models. Note that the empirical correlations are obtained after removing T_0 and T_1 from the data. In this case, the observed negative correlation at long distance is likely due to the limited domain of the simulation; as the residuals average to zero over space, if the residual is above zero in one area, it must be below zero in another. This effect is not captured by the form chosen for the covariance function, since it is likely only an outcome of the size of the simulation domain.

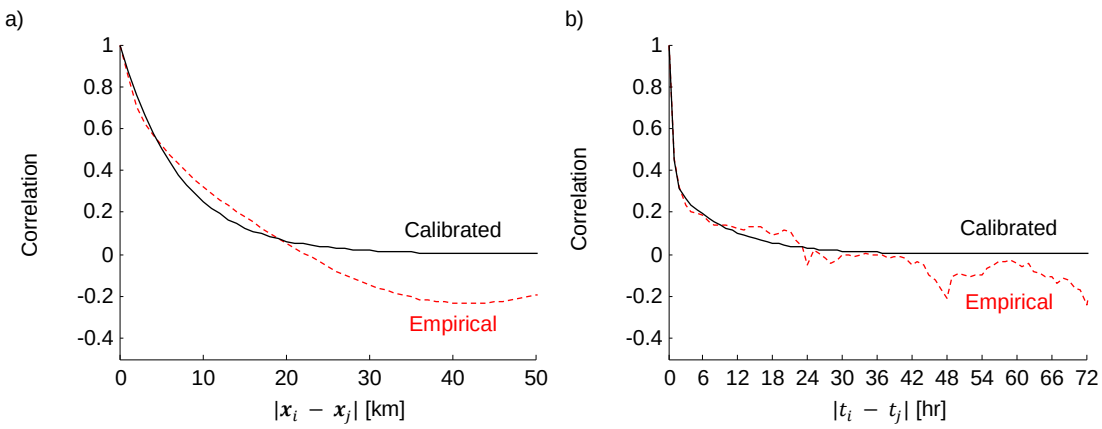

Figure 10-6: Comparisons between the empirical spatial (a) and temporal (b) correlations in the residual temperature fields for New York City, NY (dotted line), compared to the trained covariance functions of Equations 10-13 and 10-14 (solid line).

Calibrated model parameters for both New York City and Pittsburgh are listed in Table 10-D.

Table 10-D: Calibrated parameters of the covariance function for the residual temperature field.

Parameter		Value (New York)	Value (Pittsburgh)
$\lambda_{T'',X}$	[km]	7.25	7.16
$\alpha_{T'',\tau}$	[-]	0.64	0.55
$\lambda_{T'',\tau,1}$	[hr]	0.6	0.45
$\lambda_{T'',\tau,2}$	[hr]	9.5	4.1

The standard deviation of the residual field is also calibrated from the training data as the mean standard error:

$$\sigma_{T''}(\mathbf{x}) = \sqrt{\mathrm{Var}\left[\mathrm{T}(\mathbf{x},t) - \mathrm{y}_{T_0}(t) - \mathrm{y}_{T_1}(\mathbf{x},t)\right]} \tag{10-20}$$

That is, the standard deviation at location $\mathbf{x}$ is estimated as the mean standard error of the sample residual temperatures $\mathrm{T}(\mathbf{x},t) - \mathrm{y}_{T_0}(t) - \mathrm{y}_{T_1}(\mathbf{x},t)$ over the time duration of the training simulation data. Note that for this calibration step, local measurement errors $\epsilon_T(\mathbf{x},t)$ are ignored as they are assumed to be negligible compared to the variance in T''.

Figure 10-7a shows the resulting standard deviation field calibrated from training data for New York City. Note that the highest variability in temperature occurs to the northwest of the city, in the area of Orange County, New York (the green area to the upper left of Figure 10-7a). This area is more mountainous than others in the region considered, which may contribute to this variability. Coastal areas of New Jersey, Long Island, and Connecticut also appear to have more variable temperatures, as compared with the suburban areas surrounding the city.

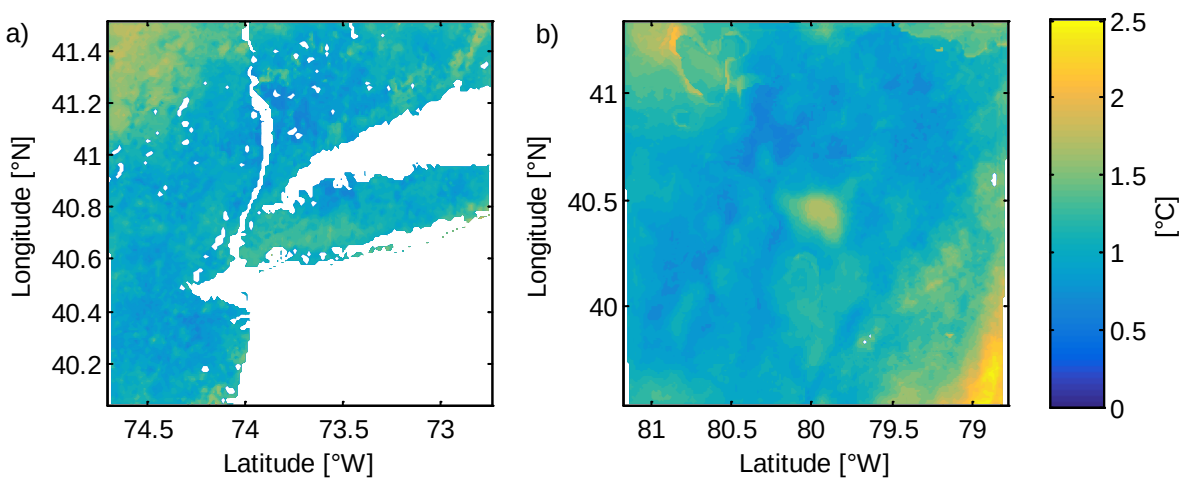

Figure 10-7: Spatial standard deviation in the residual temperature field for the region of New York City, NY (a) and Pittsburgh, PA (b).

Figure 10-7b shows the spatial distribution of the standard deviation in the residual temperature for the region around Pittsburgh. Again, this deviation is higher in the Appalachian Mountains to the southeast (the ridges to the lower right of Figure 10-1b) and to the northwest in the vicinity of Youngstown, PA (the urban area to the upper left of Figure 10-1b).

10.2.4. Validation

In this section, an assessment is made of the performance of the Gaussian process model for temperature prediction. It is assumed that the data simulated by WRF-PUCM for New York City and Pittsburgh during the validation period represent the true surface temperature field for those areas. The Gaussian process temperature models calibrated using the training data are tested on their predictive performance during the validation period (excluding any overlap with the training period) for each city. Model testing is performed under four representative use cases for the model, as discussed below.

Case 1: Local temperature measurements $\mathrm{y}_T(\mathbf{x}, t)$, as in Equation 10-19, are used to update the prior model that resulted from the calibration process. For this purpose, a random subset of the discrete locations in the region is selected to represent the locations of simulated weather stations, i.e., locations where temperature measurements are obtained and used to update the model. These temperature measurements are extracted from the WRF-PUCM simulation results during the validation period for the selected simulated weather station locations. A small error of $\sigma_\epsilon = 0.02℃$ is assumed for these measures, i.e., the local measures are assumed to be very accurate. Data collected at all of these stations prior to the current time constitute the measurement set Y, and from this information the posterior temperature prediction $\mathbf{T}|\mathbf{y}$ is obtained. Unless otherwise indicated, the random subset of locations selected to serve as simulated weather stations represents 5% of the discrete locations of the modeled domain, corresponding to an average coverage density of one station per $20\ \mathrm{km}^2$. Note that, under this case, the full training data set simulated by WRF-PUCM is only used to calibrate the model parameters; after calibration is complete, only a subset of this data corresponding to the simulated weather stations is used to condition the posterior Gaussian process model.

Case 2: In addition to the local measurements from case 1, global average measurements $\mathrm{y}_{T_0}(t)$, as in Equation 10-17, are available for past and future times. Negligible error is assumed in these global measurements. This case corresponds with the availability of an accurate deterministic predictive model for temperature, e.g., a coarse resolution WRF model for the region operated by the NWS, which provides input into the probabilistic model. This global average temperature data, together with the local measurements, constitute the measurement set Y, and from this information the posterior temperature prediction $\mathbf{T}|\mathbf{y}$ is obtained.

Case 3: This case includes global average measurements for both past and future, as in case 2, but does not include the local measurements of case 1. Instead, measurements of the cyclic temperature patterns from the training period are included as observations $\mathrm{y}_{T_1}(\mathbf{x}, t)$, as in Equation 10-18. Negligible error is assumed for these pattern measurements. This case corresponds to a scenario where prior runs of a fine resolution model are used to calibrate the cyclic temperature patterns for the region, and are combined with real-time runs of a coarse resolution predictive weather model that provides the city-wide average temperature.

Case 4: This is the same as case 3, but also includes the local measurements of case 1.

Besides these use cases, two baseline cases are also considered that do not make use of the Gaussian process framework. These baselines are provided as examples of methods for temperature prediction that do not make use of the statistical properties of the surface temperature field as encoded in the Gaussian process model. These baselines are:

Baseline 1: The global mean temperature $\mathrm{T}_0(t)$ at each time is used as an estimate of the temperature at each location. This is the same as the use of global average measurements $\mathrm{y}_{T_0}(t)$ for both past and future times, since negligible error is assumed in these forecasts. This is analogous to current weather forecasting methods where predictions are made for a large spatial region and applied for every location within the region.

Baseline 2: The temperature at any location is assumed to be the temperature measured by the nearest weather station. The weather stations in this baseline are the same simulated weather stations introduced in case 1. This corresponds to real-life prediction practices in which current temperature readings at the nearest weather station are used as a predicted temperature for the surrounding area and for the near future. In other words, this baseline case makes use of the same input information as test case 1, but does not use the Gaussian model to process this information.

Table 10-E summarizes these cases and baselines according to what inputs are provided to the Gaussian process model in the test cases or what data are used for the baseline cases. The test cases listed here are assumed to represent realistic use cases of the model under varying levels of information availability, ranging from an "information-poor" scenario with a few local weather stations to an "information-rich" scenario with fine resolution training data, accurate forecasts of the global average temperature from a coarse resolution model, and data from local weather stations. For the baseline cases, two representative scenarios are also considered, corresponding to the use of forecasts from a coarse resolution weather model or to measurements from the nearest weather station. These are considered to be typical, intuitive methods for temperature prediction in urban areas.

Table 10-E: Test cases for model validation.

Case	$\mathrm{y}_{T_0}(t)$ (accurate forecast)	$\mathrm{y}_{T_1}(\mathbf{x},t)$ (fine resolution training data)	$\mathrm{y}_{T}(\mathbf{x},t)$ (local measurements)
1			√
2	√		√
3	√	√	
4	√	√	√
BL-1	√		
BL-2			√

10.2.4.1. New York City

Validation is conducted on a reduced domain of the full model, a 30 by 30 km area around Manhattan Island, in order to reduce the computational cost for updating the posterior model for the entire domain (a joint Gaussian distribution over more than 25 thousand variables at each time step) to a more manageable size (nearly 1 thousand variables at each time step). For all cases, 5 days of past local measurement data are used to condition the posterior model, i.e., 5 days of past data are used for the local measurements $y_T(\mathbf{x}, t)$, while measures $y_{T_1}(\mathbf{x}, t)$ are derived from the training data set only (as the cyclic patterns are assumed not to change over time), and measures $y_{T_0}(t)$ are only needed for the prediction time, since these represent accurate forecasts from an external source. This allows the amount of data considered to be reduced, and improves the efficiency of the model updating. Truncation of data records is an intuitive solution to reducing the computational complexity of the updating problem in this application, since more recent temperature measurements will largely "override" these older measurements during the updating process. For comparison, while the deterministic WRF-PUCM model takes about 12 hours to generate one week of simulated temperatures (at the spatial and temporal resolutions listed in Table 10-A), the calibrated probabilistic model can be run in two minutes (in Case 3) to two hours (in Case 4) to obtain posterior distributions for a week's worth of temperature fields.

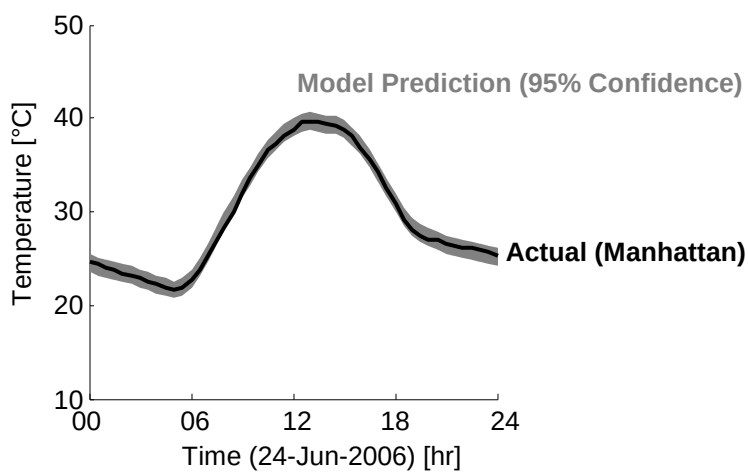

Figure 10-8: Comparison of the temperature record from WRF (black line) for a location on Manhattan Island with the 95% confidence interval (grey area) of the posterior model for one-hour-ahead prediction under use case 4.

As an example of the output of the model, Figure 10-8 compares the WRF-simulated (actual) temperature record at a location on Manhattan Island for a representative day during the validation period with the 95% confidence interval of the posterior model at that location. The model is updated as described in use case 4, with the nearest local measurement station being 4 km away. One-hour-ahead predictions are shown in this figure (i.e., the posterior model is updated using local measurements taken prior to one hour before the time at which temperature

is to be predicted). This demonstrates how the model can provide probabilistic temperature estimates for specific locations ahead of time, which quantify the uncertainty in the prediction.

The predictive performance of the calibrated model is evaluated on the validation data set. Prediction is performed using the calibrated prior model, together with a certain set of temperature observations, as per the use case, to obtain a posterior model for future temperatures. The posterior mean field of the Gaussian process is then used as the temperature field prediction. Predictive accuracy is assessed based on the average root mean square error (RMSE) between the model's temperature predictions and the simulated temperature values in the validation data across the domain being modeled. The validation data set is divided into 10 subsets, prediction is performed under each case for each subset, and the RMSE is averaged across these repetitions. Figure 10-9 indicates the results of this assessment. Table 10-F provides a summary of the predictive performance of the model for certain lead times.

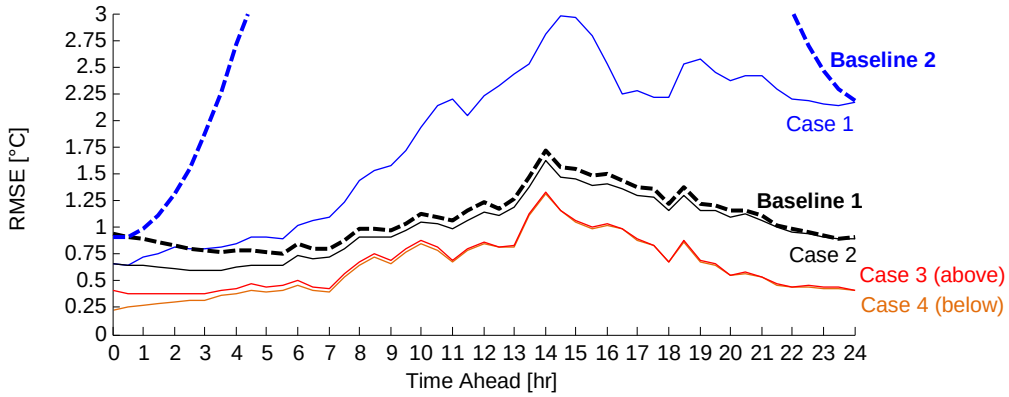

Figure 10-9: Predictive performance of the probabilistic model for New York City, NY, under various use cases. For comparison, the performances of two baseline prediction cases are also provided.

In terms of the test cases, predictive performance improves between cases 1, 2, 3, and 4, as more information is included for updating the model, as is expected. For the base cases, performance is comparable in the short-term, with zero-lead-time performance of the second baseline slightly exceeding that of the first. However, performance of the second baseline is much poorer than other models as the forecasting horizon increases and the most recent measures become increasingly out-of-date. The first test case, using the same input information as the second baseline, performs fairly well in the short-term, but performance also declines with increasing forecasting horizon as this information becomes out-of-date. Performance improves slightly at about 24 hours prediction lead time, due to cyclic correlation in the global average and temperature patterns. For case 2, accurate forecasts of the global average temperature improve long-term performance greatly. In the short-term, performance of the second test case exceeds that of the first baseline, due to the inclusion of local measurement data. However, as time progresses and these local data become outdated, performance of the second test case is

comparable to that of the first baseline. The third and fourth cases, due to the inclusion of training data for spatial temperature patterns, outperform the first baseline, with the fourth case outperforming the third in the short-term due to incorporation of the local data.

For comparison, the RMSE of the WRF-PUCM model with respect to measurements at NWS stations in the simulation domain has been approximated as 2°C for near-surface air temperatures (Ramamurthy et al., 2015). The performance of the posterior model compares favorably with this value across all cases and lead times considered here. Note that, although the fitting between the probabilistic model and the simulated dataset is closer than that between the dataset and measured real-world temperatures (indicating potential overfitting of the model to the data), the probabilistic approach can be applied generally to other deterministic models and to other temperature data sets. As more accurate deterministic modeling methods and data from dense urban temperature monitoring networks become available, a probabilistic model calibrated using this information should be able to provide comparable predictive performance. Since even a 1ºC difference in surface temperatures can result in a 4.5% change in heat mortality risk (Anderson and Bell, 2011), accurate prediction of these temperatures, such as provided by the calibrated Gaussian process model conditioned on local temperature measurements, allows for a better understanding of the impacts of extreme heat on urban residents.

Table 10-F: Comparative results for New York City, listing RMSE for different model cases and prediction lead times.

Method	**RMSE [°C] of Lead Time [hr]**			
	0	1	6	24
Case 1	0.66	0.72	1.02	2.17
Case 2	0.66	0.64	0.74	0.90
Case 3	0.42	0.38	0.51	0.42
Case 4	0.23	0.27	0.46	0.41
Baseline 1	0.94	0.89	0.84	0.91
Baseline 2	0.90	0.99	4.49	2.18

The predictive performance of the model is also assessed from a probabilistic standpoint. For this purpose, a representative subset of the prediction errors discussed above are plotted in a quantile-quantile plot in Figure 10-10, comparing their distribution against a Gaussian distribution. Ideally, if the temperature was perfectly described by a Gaussian distribution, these errors would lie along a straight line. This figure indicates that the errors have a nearly Gaussian distribution for about two standard deviations above and below the mean, covering the middle 95% of the data under all use cases. However, as the true temperature field is non-Gaussian, the tails of the distributions do not exactly match. In fact, Figure 10-10 indicates that the Gaussian model tends to over-predict positive errors and under-predict negative ones, i.e., the model predicts a wider distribution for temperatures than is actually observed. This phenomenon is more severe under use cases 3 and 4 than under cases 1 or 2. For instance, for extremely high temperatures under

use case 4, events with an exceedance probability in the data of 10^{-5} will be predicted by the model to have an exceedance probability of 10^{-3}. This shortcoming of the model can be addressed using a nonlinear transformation function between the Gaussian process model and the surface temperature. However, for the purposes of this work, this approach is not adopted, and the resulting inaccuracies are taken as an acceptable source of error in the presented results.

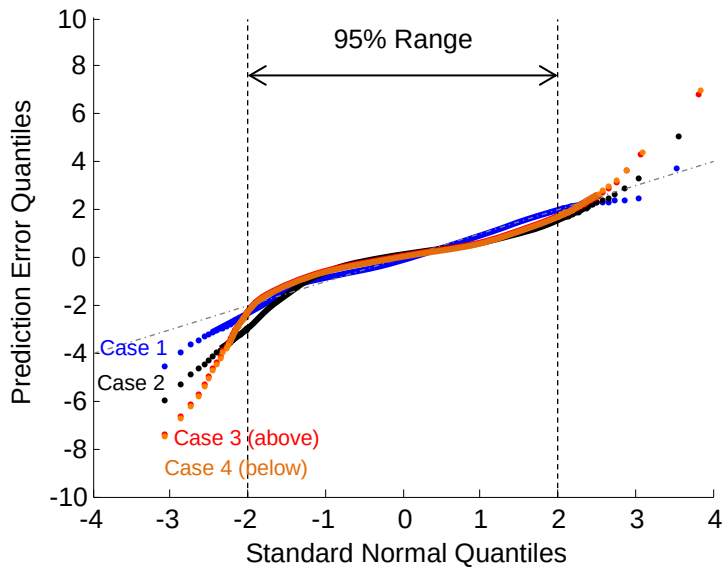

Figure 10-10: Quantile-quantile plot of 1,000 empirical residuals of the probabilistic model for New York City, NY, under each use case. The 95% confidence interval of the Gaussian distribution is indicated between the vertical dashed lines. The grey dashed line indicates the ideal normal distribution.

Finally, Figure 10-11 reports the performance of the posterior model, under use case 4 and for no prediction lead time, using various percentages (0%, 2%, 5%, 10%, 20%, 50%, and 100%) of discrete locations as simulated weather stations providing input for model updating. Intuitively, increasing the available data improves the model predictions, although the marginal benefit of adding data is higher when the total amount of data is lower, as indicated by the positive curvature of the left part of the graph. With 100% of locations being sampled, the prediction accuracy is limited only by the noise of these measurements, $\epsilon_T(\mathbf{x}, t)$, in this case assumed to be 0.02°C.

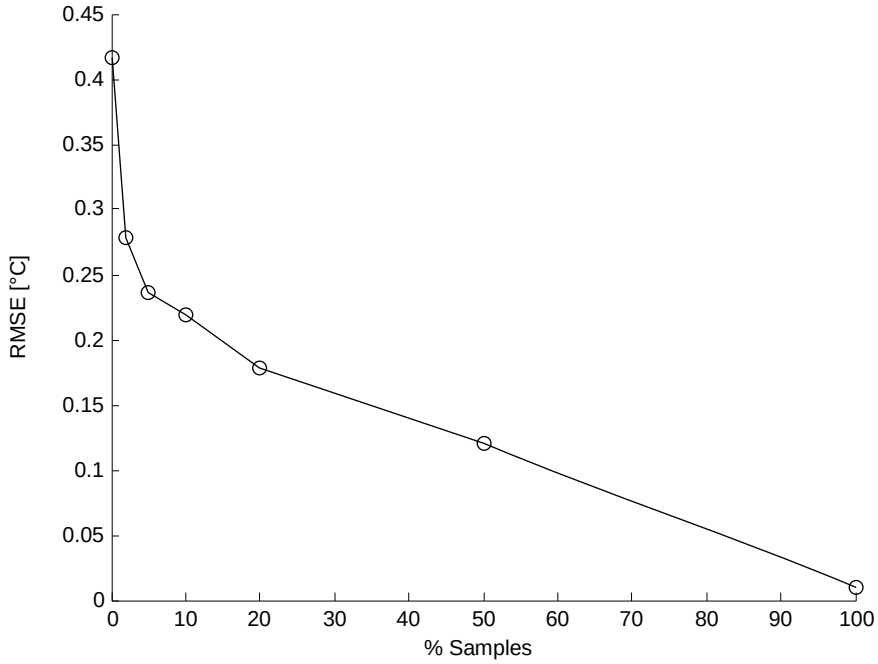

Figure 10-11: Results for posterior model performance under use case 4 with different levels of data (as measured by % Samples, or the fraction of discrete locations at which data are collected for updating the model) for New York City, NY.

10.2.4.2. Pittsburgh

An evaluation of the performance of the calibrated model on the validation set is also conducted for Pittsburgh. Analysis is restricted to the city proper, an area of 122 square kilometers within the simulation region. Table 10-G reports results of the validation, which are mostly qualitatively similar to results for New York City. Interestingly, model performance in case 1 for zero lead time slightly exceeds its performance in case 4, and in the second baseline case also, performance is comparable to the model in case 4. This is likely the result of the relatively small size of the region under consideration in the model; measures at only a small number of locations within the city serve as an adequate representation of the temperature field, even without the application of the Gaussian Process model. However, the utility of this model is clearer for greater lead times, as the performance of the model improves with increasing data for nonzero lead times and these models show improved performance over comparable baseline cases, as previously noted for the example of New York City.

Table 10-G: Comparative results in Pittsburgh, listing RMSE for different model cases and prediction lead times.

Method	**RMSE [°C] of Lead Time [hr]**			
	0	1	6	24
Case 1	0.25	0.78	2.07	3.51
Case 2	0.38	0.73	1.29	1.12
Case 3	0.60	0.57	0.50	0.42
Case 4	0.27	0.43	0.51	0.42
Baseline 1	1.54	1.46	1.10	1.73
Baseline 2	0.27	1.39	6.71	3.41

Finally, a demonstration is provided of how the proposed methodology can account for the UHI effect. Simulated temperatures from the validation data set at a representative location in downtown Pittsburgh are compared to the temperatures at the location of the nearest NWS station to Pittsburgh, located at Pittsburgh International Airport (PIT), about 20 km distant from the downtown area. In this comparison, the simulated temperature data at the location of PIT, as would be measured by the weather station, are used as an estimate of the temperature in the downtown area. This is an example of the second baseline case, with only a single measurement location. Alternatively, these data can be fed into the Gaussian process temperature model according to use case 4 (again with only the single local measurement location at PIT) to obtain a posterior prediction for temperature at the representative downtown location. One-hour-ahead predictions are considered, with the baseline case using the most recent measure of temperature at PIT as the one-hour-ahead prediction. Figure 10-12a indicates the results of this comparison for a representative day during the validation period.

The downtown temperature tends to be higher than that measured at the airport outside the city due to the UHI effect (although other factors such as elevation differences may contribute to the difference too). The measurements taken at PIT underestimate the downtown temperature by 1.5°C on average during the validation period. The RMSE for the prediction is 2.3°C; note that this is the prediction error for a specific point, rather than an average over the region as in Table 10-G. By contrast, the posterior mean of the Gaussian process model under use case 4 with local measurements taken at PIT overestimates the downtown temperature at the representative location by only 0.1°C on average during the validation period, with a RMSE of 0.7°C. This demonstrates how the Gaussian process model can be used to account for systematic temperature differences in an urban area due to the UHI effect, using limited local measurement data together with prior fine resolution simulation data for the region in question as well as accurate regional temperature forecasts.

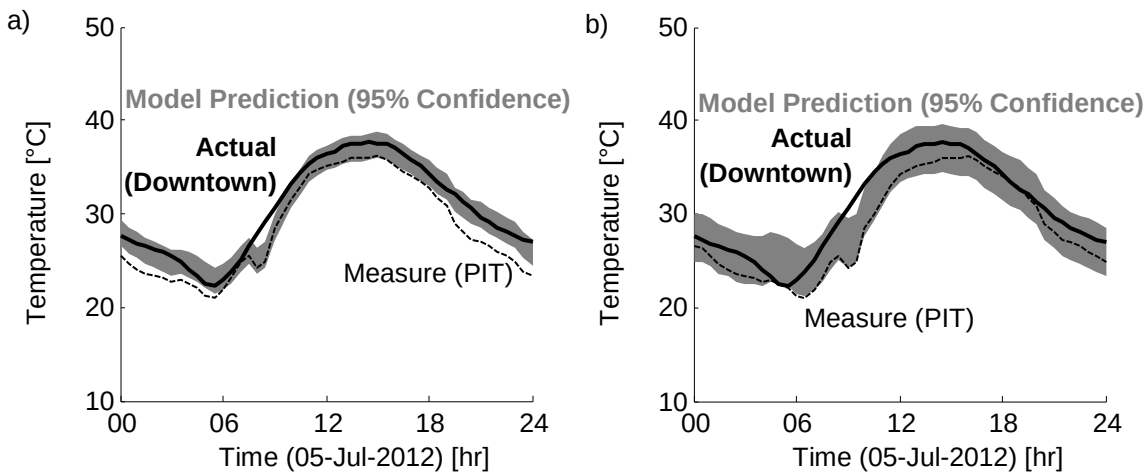

Figure 10-12: Comparison of actual temperature in downtown Pittsburgh from the simulated temperature data of the validation period on the date indicated (solid line) with the measurement of the nearest weather station (dotted line) and the 95% confidence interval of the posterior prediction of the Gaussian process temperature model (grey area) for one-hour-ahead prediction under case 4 with local data from the single weather station (a) and using these local measurement data as well as prior measures of the cyclic pattern obtained from training data (b).

Furthermore, the performance of the model is assessed for predicting the local temperature at a representative location without accurate regional forecasting, i.e., only using temperature measurements taken at PIT and samples of the cyclic temperature pattern obtained from the training data. Note that this does not correspond to one of the cases listed in Table 10-E, but rather is a demonstrative case combining local measurement data $y_T(\mathbf{x}, t)$ obtained at a single point with prior measures of the temperature pattern $y_{T_1}(\mathbf{x}, t)$. The performance of the model for one-hour-ahead prediction is again compared with that of the second baseline case using the same local measurement data. Results for the representative day are indicated in Figure 10-12b. The PIT measures still underestimate downtown temperatures by 1.5°C on average, with a RMSE of 2.3°C, while the posterior model overestimates temperatures by 0.2°C, with a RMSE of 1.6°C. This demonstrates how the modeling framework (using the same locally measured temperature data together with prior fine resolution simulations of the urban temperature pattern) can account for systematic differences in temperature such as those due to the UHI effect to reduce prediction bias with respect to the local measurement data alone.

10.3. Temperature Monitoring

Using the probabilistic model for urban temperatures developed in Section 10.2, temperature monitoring can be performed using the metrics for sensor placement discussed previously in this work. In Section 10.3.1, the probabilistic temperature model is combined with a model for urban temperature vulnerability to characterize urban heat risk probabilistically. In Section 10.3.2,

specific formulations for the conditional entropy, weighted prediction error, and VoI metrics in the context of urban temperature prediction and decision-making are presented.

10.3.1. Risk Formulation

In this chapter, risk is quantified as the product of a factor relating to the hazard source (in this case, the urban temperature, as described by the probabilistic model developed in Section 10.2) and a factor relating to the vulnerability and exposure of the urban population to this hazard. Heat vulnerability is quantified via a heat vulnerability index developed by Bradford et al. (2015) for the city of Pittsburgh. This index was created after analyzing 13 socio-economic metrics, such as age, income, and education, which are identified in the literature as contributing to heat-related vulnerability, and conducting a factor analysis to group these metrics into factors that explain at least 70% of the variance in all metrics (after removing metrics that failed a multicollinearity test). The vulnerability index was then created by binning these factors into 6 increments for each census block of the city, and averaging across all factors. The result is an index for each census block between 1 and 6, with 6 corresponding to a higher vulnerability (with the factors being on average 2 or more standard deviations above the mean) and 1 to a lower vulnerability (factors 2 or more standard deviations below the mean). The results of this study are translated into a vulnerability index field, $\mathrm{V}(\mathbf{x})$, for the city of Pittsburgh. This is done by taking the area-weighted averages of vulnerability indices for census blocks within the area associated with spatial coordinate $\mathbf{x}$ (a one-square-kilometer box around the coordinate) and assigning this weighted average as the vulnerability $\mathrm{V}(\mathbf{x})$ for that location. For reference, this vulnerability field is depicted in Figure 10-13. This vulnerability is assumed to be constant with respect to time; variations in population distributions throughout the day are not considered in this vulnerability index.

Exposure is not considered separately from vulnerability in this work. In principle, a spatio-temporal model of urban population movements might be created to describe the varying distribution of people in a city throughout the day, and potentially how this is altered by extreme temperatures. However, no such models were identified in the literature, and the development of a separate model is beyond the scope of this work. Furthermore, the vulnerability index used here already considers exposure-related factors such as access to air conditioning and to green spaces. Therefore, exposure is considered as a binary value that is either present or absent, following the approach of Grower et al. (2011), and it is assumed that all areas except those dominated by water have populations exposed to extreme heat.

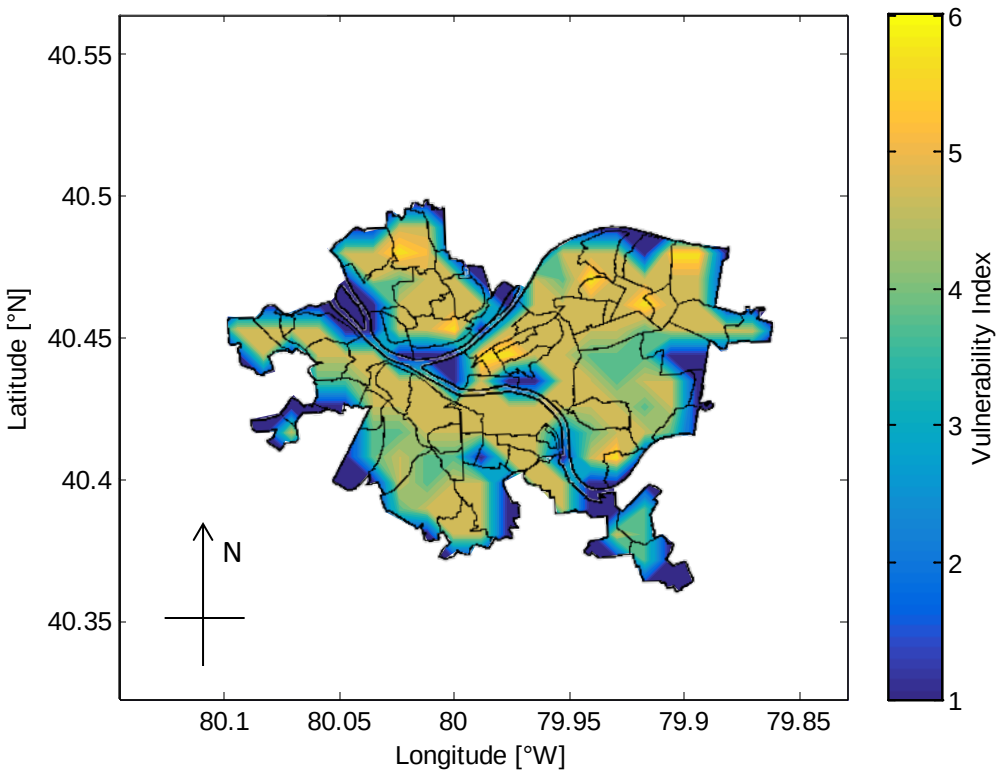

Figure 10-13: Vulnerability index map for Pittsburgh, PA.

Combining the hazard value, i.e., the spatio-temporal surface temperature $\mathrm{T}(\mathbf{x}, t)$, with the spatial vulnerability index $\mathrm{V}(\mathbf{x})$ defines a spatio-temporal risk:

$$\mathrm{R}(\mathbf{x}, t) = \mathrm{T}(\mathbf{x}, t)\, \mathrm{V}(\mathbf{x}) \tag{10-21}$$

As the product of a Gaussian process with a deterministic scaling field, the spatio-temporal risk is itself described by a Gaussian process model.

The probabilistic character of this model allows for other commonly used heat risk indices to be derived from it. For example, previous studies (e.g., Buscail et al., 2012; Harlan et al., 2012) have developed heat risk indices based on temperatures observed via satellite at a specific time during a heat wave event. Such information can be integrated into the Gaussian process framework to define a posterior risk metric conditioned on satellite temperature data. Alternatively, Aubrecht and Özceylan (2013) define a risk metric using the number of heat wave days (i.e., days with peak temperatures exceeding a given threshold) in a given year as a hazard measure. Such a metric can be obtained from the probabilistic heat risk model above, e.g., by conducting a Monte Carlo simulation from the probabilistic temperature model and counting the number of days the threshold is exceeded. Furthermore, risk based on more complex thermal comfort indices (e.g., Höppe, 1999) can also be evaluated.

To illustrate how the probabilistic model described above can be used to generate a heat risk index map (such as those of previous work on heat risk quantification), an example is presented below. For this example, risk is defined based on the 95th percentile temperature at the hottest time of the day for a given location. Denote by $\mathrm{T}_{95\%}(\mathbf{x}, t)$ the 95th percentile temperature at location $\mathbf{x}$ and time t, predicted by the probabilistic temperature model, and take the maximum of these temperatures over a typical 24 hour period as the hazard measure. Note that this is only one method for describing the hazard, and is used here for illustrative purposes only. Finally, to allow comparison between this hazard measure and the vulnerability index, these maximum 95th percentile temperatures are binned into 6 increments based on the range of these values across the city. This is denoted by $\mathrm{bin}[\cdot]$, which describes a mapping from the continuous temperature values to an integer between 1 and 6 (this is the same binning function used to develop the vulnerability index, as described above). Following this transformation, the hazard and vulnerability measures are both indices ranging from 1 to 6, and can be directly multiplied to define a risk metric:

$$\mathrm{R}_{95\%}(\mathbf{x}) = \mathrm{bin}\left[\max_{t\in[0,24]} \mathrm{T}_{95\%}(\mathbf{x}, t)\right] \mathrm{V}(\mathbf{x}) \tag{10-22}$$

The model $\mathrm{T}(\mathbf{x}, t)$ used to define these percentiles is calibrated for the region of Pittsburgh and conditioned on measurements $\mathrm{y}_{T_1}(\mathbf{x}, t)$ obtained from the training data set. The resulting risk index map is shown in Figure 10-14, with the risk metric normalized to vary between 0 and 1. High-risk areas are mostly located to the center of the region, i.e., the downtown Pittsburgh area, as well as to the north, both to the northeast and northwest. While the southwest of Pittsburgh has comparable vulnerability other areas, as shown in Figure 10-13, lower extreme temperatures here lead to lower risk indices in this area.

Many heat risk index mappings, such as that of Equation 10-22 involve non-linear functions of the temperature field, and so are no longer Gaussian. Therefore, it may be difficult to analyze these mappings directly in order to determine the utility of sensor placements in reducing uncertainty or in supporting risk mitigation actions. Furthermore, the use of such static mappings ignores the changing nature of the hazard and risk over time, as well as over space, the effects of which are lost when only a single arbitrarily-selected feature (e.g., the 95th percentile temperature) of the full spatio-temporal hazard model is used. For these reasons, the following sections apply sensor placement metrics directly to the underlying Gaussian temperature or risk fields (as in Equation 10-21). The weighted prediction error metric, discussed in Section 10.3.2.2, is in fact a metric measuring uncertainty in this risk field. Sensor placements resulting from these various objectives are compared to the risk field using the index of Equation 10-22 in Section 10.4.1.4.

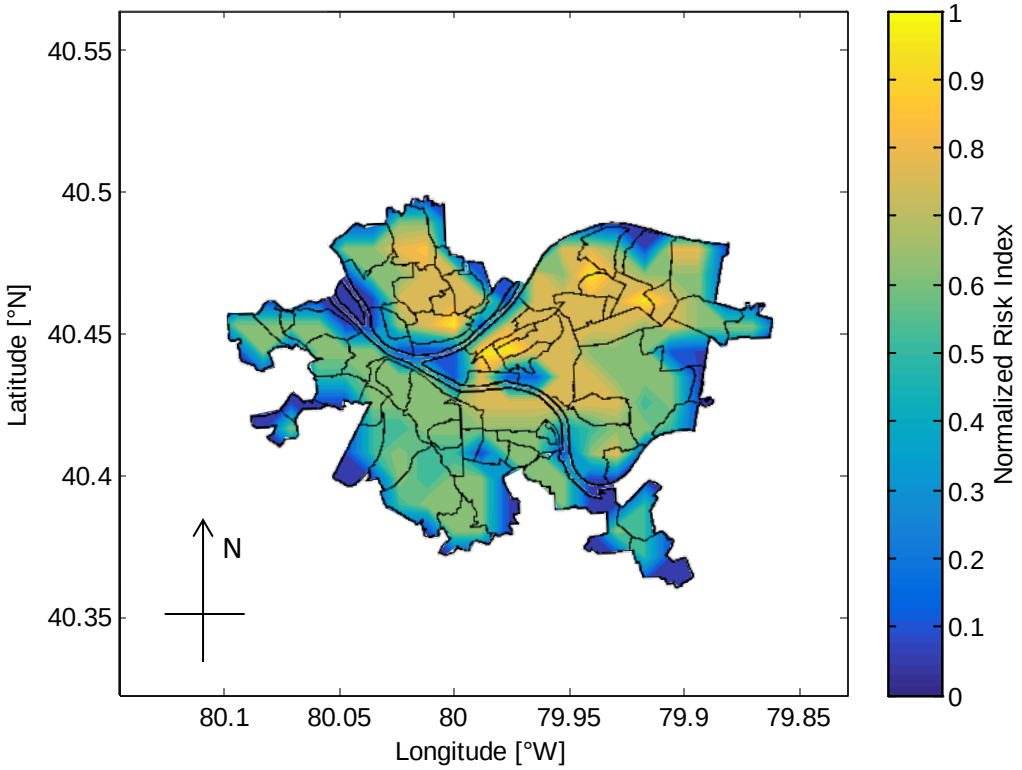

Figure 10-14: An example risk index based on 95th percentile peak daily temperature in Pittsburgh, PA, and the vulnerability map of Figure 10-13. This index is normalized between 0 and 1.

10.3.2. Sensing Metrics Considered

Many factors should be considered when selecting locations to measure temperatures in an urban area. The uncertainty in the temperature field, the vulnerability of people in different areas to extreme heat, and the available actions (and their consequences) that will be undertaken in response to the gathered information are all contributing factors to the temperature sensing problem. An understanding of the spatially distributed risk, such as might be gained from Figure 10-14, is not necessarily sufficient. For example, monitoring temperatures where heat risk is high is an intuitive solution, but may neglect areas where additional information could reduce risk uncertainty; conversely, focusing on areas of high risk uncertainty but low expected value is also inefficient. In this section, three metrics for sensor placement applied to the problem of urban temperature monitoring are described: a conditional-entropy-based objective, an objective based on a weighted average predictive performance of the temperature model, and a VoI-based objective. Each objective corresponds to a particular viewpoint on what information is important to temperature monitoring. Under the conditional entropy metric, because the contribution of vulnerability to the risk is already known, reduction of uncertainty in the temperature field is the

goal of sensing. Under the weighted prediction error metric, because risk is proportional to the product of temperature and vulnerability, uncertainty in this product is to be reduced. VoI uses a hazard-based decision problem with utility of decision outcomes proportional to vulnerability to identify which additional information might best support risk response decision-making.

10.3.2.1. Conditional entropy of the temperature field

The conditional entropy of the random temperature field is used to define one objective for sensor placement:

$$\mathrm{M}_{\mathrm{Ent}}(Y) = \mathrm{H}(T) - \mathrm{H}(T|Y) \tag{10-23}$$

where $\mathrm{H}(T|Y)$ denotes the posterior entropy of the temperature field T conditioned on measurement set Y. This entropy can be evaluated for the prior and posterior Gaussian temperature fields as in Equation 4-44. Note that, rather than the cost-benefit formulation of Equation 4-41, a cost-neutral approach is used here, i.e., the cost of measurements is not considered. This objective is optimized using the general optimization objective of Equation 2-16, together with the greedy solution approach of Section 2.2.1.

10.3.2.2. Weighted prediction error of the temperature field

A weighted prediction error metric, such as described in Section 4.4.1.1, is also used. Prediction error is measured by the sum of square differences between the actual temperature and predicted temperature at a specific time and location. Furthermore, these prediction errors are weighted by the local vulnerability index; the weight matrix $\mathbf{Z}$ of Equation 4-35 is defined using the vulnerability field $\mathrm{V}(\mathbf{x})$, with vulnerability being constant over time. In this way, prediction errors in vulnerable areas are penalized more heavily by the metric, with penalties proportional to the product of the local vulnerability and the posterior variance in temperature. Under this choice of weighting, the weighted prediction error metric provides a measure of the uncertainty in the risk field $\mathrm{R}(\mathbf{x}, t)$ of Equation 10-21.

For defining an objective based on this weighted prediction error loss function, sensing costs are again ignored, as with conditional entropy above. The resulting sensor placement objective is:

$$\mathrm{M}_{\mathrm{Err}}(Y) = \mathbb{E}\mathrm{L}_{\mathrm{pred}}(\emptyset) - \mathbb{E}\mathrm{L}_{\mathrm{pred}}(Y) \tag{10-24}$$

This objective is computed following the approach outlined in Section 5.2.1 and is greedily optimized following Section 2.2.1.

10.3.2.3. Value of information for heat wave advisory issuance

To define the VoI metric, the vulnerability index is used to define a loss function for a decision-making problem motivated by heat advisory issuance. Consider that a cost $\mathrm{C}_f(\mathbf{x})$ is incurred in

an area if the temperature exceeds a certain threshold T_{limit}. This cost captures the consequences of the area's population being exposed to extreme temperatures. The city government has the option of issuing an advisory for this area at this time, i.e., by choosing $\mathrm{a}(\mathbf{x}, t) \in \{0,1\}$, where a choice of "1" corresponds with the issuance of a heat advisory, and a choice of "0" corresponds with no warnings being issued. By choosing to issue a warning, people in the region will be encouraged to seek shelter, and the consequences of exposure will be avoided. There is also a cost for issuing this advisory, $\mathrm{C}_r(\mathbf{x})$, which quantifies the effort necessary to issue the warning and the loss of productivity in the area where the advisory is issued due to restrictions on outdoor activity.

This decision-making problem defines a spatio-temporally decomposable loss function as in Equation 7-7 for the discretized temperature field. The local spatio-temporal loss is expressed following the prototypical loss function of Equation 4-15 as:

$$\mathrm{L}_i^{(t)}\left(\mathbf{s}_i^{(t)}, \mathbf{a}_i^{(t)}\right) = \mathrm{C}_f(\mathbf{x}_i)\left(1 - \mathbf{s}_i^{(t)}\right)\left(1 - \mathbf{a}_i^{(t)}\right) + \mathrm{C}_r(\mathbf{x}_i)\mathbf{a}_i^{(t)} \tag{10-25}$$

where $\mathbf{a}_i^{(t)} = \mathrm{a}(\mathbf{x}_i, t)$ and $\mathbf{s}_i^{(t)}$ is defined following Equation 3-14, with the limit-state function defined following Equation 3-11 with

$$\mathbf{g}_i^{(t)} = T_{limit} - \mathrm{T}(\mathbf{x}_i, t) \tag{10-26}$$

That is, the state is defined by whether or not the local temperature exceeds the given threshold at a given time. Because of the relatively short time horizon of the problem, no discounting is assumed; instead a factor of $\gamma(t) = \frac{1}{n_\tau}$ is used. The VoI metric defined following this loss function represents the average hourly reduction in expected loss in this heat advisory issuance decision-making problem over the domain of the city. In this case, sensing costs are considered, and the VoI objective is defined as in Equation 4-12.

Ideally, both $\mathrm{C}_f(\mathbf{x})$ and $\mathrm{C}_r(\mathbf{x})$ can be quantified in common units, e.g., by converting mortality, morbidity, and productivity loss to common monetary quantities using accepted practices. However, such a quantification is beyond the scope of this work. As discussed by Bradford et al. (2015), the heat vulnerability index is not a measure of the consequences of a heat wave event. It is, rather, one input to a utility function (or loss function) for decision-making in response to extreme temperatures. The relationship between the societal factors captured in this index and the impacts of extreme temperatures, e.g., in terms of hospitalizations due to heat stress, remains unclear and is an open area of research (e.g., Aubrecht and Özceylan, 2013; Hondula and Barnett, 2014; Bradford et al., 2015). Instead, it is assumed for illustrative purposes that consequences are equal to the vulnerability index of an area, and so $\mathrm{C}_f(\mathbf{x}) = \mathrm{V}(\mathbf{x})$. Using this same vulnerability scale (ranging from 1 to 6), the consequence of issuing a heat advisory is arbitrarily chosen as $\mathrm{C}_r(\mathbf{x}) = 3\ \forall\ \mathbf{x} \in X$. Under these assumptions, for some areas of the city with low vulnerability, the consequence of being exposed to extreme heat (for this example,

extreme heat is defined by the limit $T_{limit} = 30°C$) is lower than the consequence of issuing an advisory, and so it will never be beneficial to issue a heat advisory. However, for more at-risk areas, it may be of interest to issue an advisory to avoid higher consequences of extreme heat exposure.

10.4. Sensor Placement Results

Sensor placement is conducted using the trained model for urban temperatures in Pittsburgh, as developed in Section 10.2, under the three objectives discussed in Section 10.3. In Section 10.4.1, it is assumed that (i) accurate predictions of the average temperature $y_{T_0}(t)$ over the region, (ii) prior assessments of the cyclic temperature pattern $y_{T_1}(\mathbf{x}, t)$ obtained from fine resolution simulations, and (iii) local temperature measurements $y_T(\mathbf{x}, t)$ whose locations are being optimized are available and used to update the prior spatio-temporal probabilistic model. This corresponds with use case 4 of Table 10-E. This is potentially the most realistic use case for the model and its application to improving temperature field prediction because accurate regional temperature forecasts are commonly available from coarse resolution weather simulations (e.g., as provided by the NWS) and historical simulations of urban temperature patterns can be used to obtain measurements of the cyclic temperature patterns (e.g., as generated by WRF-PUCM). In Section 10.4.2, assumptions (i) and (ii) above are dropped. This corresponds to use case 1 of Table 10-E, and to a problem of sensor placement in data-poor situations, e.g., where estimates of regional temperatures and simulation results of historical temperature patterns are unavailable.

Only potential measurements at grid points within the city are considered. The set of all possible measurement locations, denoted $Y_{T,All}$, consists of 122 sensors, one for each grid point. It is further assumed that once placed, a sensor will continuously gather temperature data over time at a negligible additional cost. Thus the problem of simultaneous sensor placement and scheduling is simplified to sensor placement only for the local temperature measures within the domain. Simultaneous placement and scheduling can also be accomplished, as discussed in Chapter 7, at the cost of a linear increase in computational complexity of the greedy algorithm as the time horizon for sensor scheduling increases.

10.4.1. With Global Mean and Cyclic Pattern Measures

The following subsections present results for sensor placement of local temperature measures under the assumption that measures of the global average temperature and the cyclic temperature patterns will be available as well (i.e., under case 4 of Table 10-E).

10.4.1.1. Conditional entropy of the temperature field

First, sensor placements are optimized for the reduction of conditional entropy of the posterior temperature field over the city of Pittsburgh. Only the joint entropy of the temperature field within the city is used to evaluate the conditional entropy metric. The results indicated in Figure

10-15 plot the locations of 9 greedily optimized temperature measurements throughout the city, together with the posterior standard deviation of the residual temperature field conditioned on these measurements. Here, sensor locations are distributed throughout the city to reduce overall uncertainty in temperature values.

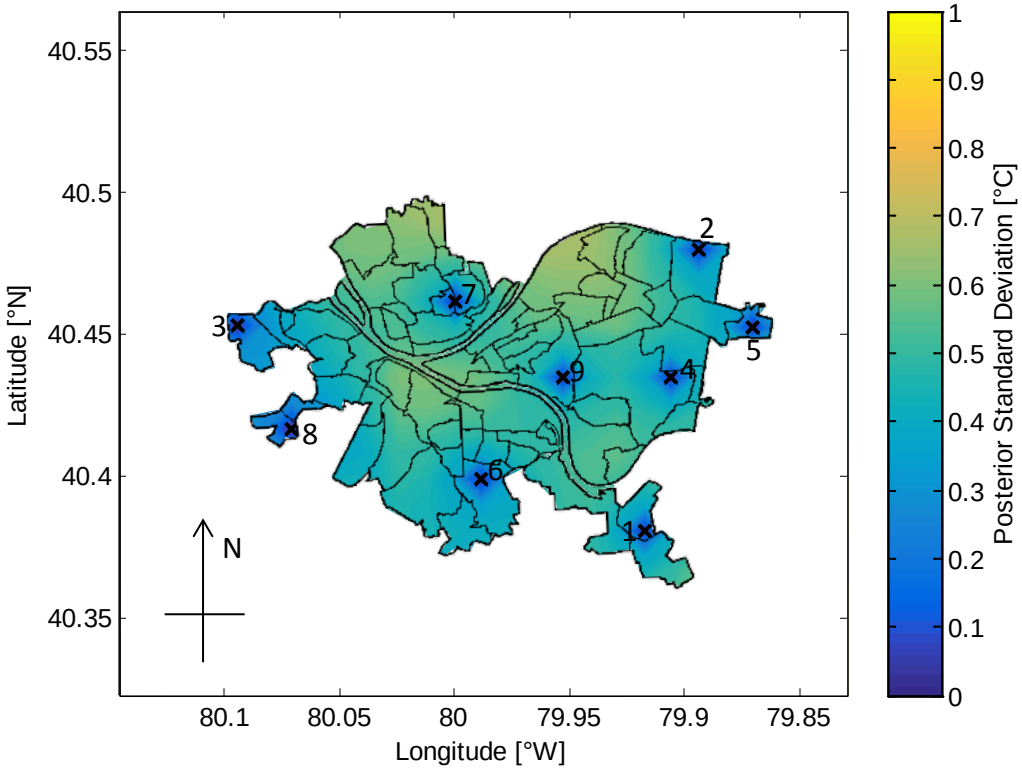

Figure 10-15: Optimized measurement locations for Pittsburgh, PA, supporting the reduction of conditional entropy of the temperature field. Sensing locations are indicated by x's, while the background indicates the posterior standard deviation of the temperature field conditional to these measurements. Numbers next to the sensing locations indicate the order the placements were selected by the greedy optimization algorithm. The city's outline is also indicated.

Note that the conditional entropy decreases as more sensors are considered, and therefore the set $Y_{T,All}$ would minimize the conditional entropy. However, the use of this set is assumed to be impractical (due to the high cost of installing so many sensors), and instead a subset of these locations is selected. The number of sensors to include in this set is chosen such that the sets optimized under each objective considered in this section have the same number of sensors, allowing for more direct comparison between them. The conditional entropy metric value, i.e., the reduction in conditional entropy as more sensors are added to the selected set, is nearly linear in the number of sensors, as shown in Figure 10-16. This is a consequence of the definition of the conditional entropy metric for multivariate Gaussian distributions as the log-determinant of the

posterior covariance matrix $\mathbf{\Sigma}_{T|Y}$. This metric is thus a logarithmic function of the probabilistic volume spanned by the distribution, quantified via the determinant of the posterior covariance matrix. As sensors are placed, this posterior covariance matrix is compressed along one dimension corresponding to the variable that is being measured. For independent variables, this entropy reduction would therefore be linear with respect to the number of sensors. Correlations within $\mathbf{\Sigma}_T$ account for the slight difference from linearity observed in Figure 10-16. This near-linearity is a general property of the conditional entropy metric applied to two-dimensional Gaussian processes. The number of sensors indicated in Figure 10-15 corresponds to the number beyond which the additional benefit provided by the next sensor is less than a given percentage of that provided by the first sensor selected. Because of the nearly linear rate of reduction, a high percentage of the initial value of 82% is chosen to identify the number of sensors to include in the set as 9, the same number selected as the optimum number by the VoI metric in Section 10.4.1.3.

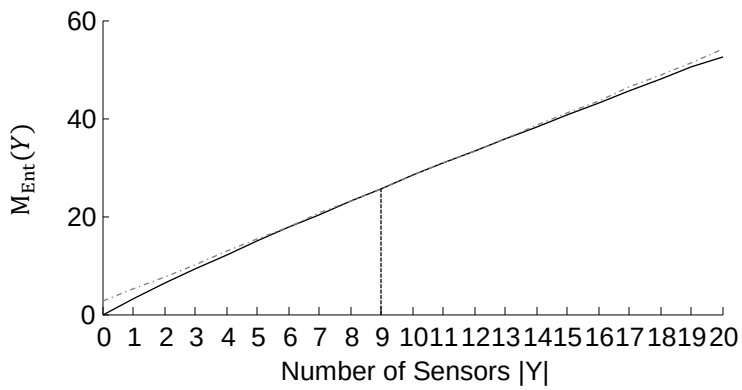

Figure 10-16: Reduction in entropy resulting from the sensor placements through greedy optimization (solid line). The line with slope corresponding to 82% of the reduction of the first sensor is also plotted (dashed line), and the point at which this line is tangent to the entropy reduction curve identifies the optimal number of sensors (in this case 9).

10.4.1.2. Weighted prediction error of the temperature field

Next, sensors are optimized based on the residual prediction error weighted by the local vulnerability to extreme temperatures. The results are indicated in Figure 10-17, plotting the locations of 9 greedily optimized sensor placements in the city that reduce the overall weighted prediction error. Because of weighting by the vulnerability index, the resulting sensing scheme is slightly different from that indicated in Figure 10-15, although one sensor placement is common to both. Here, sensor placements generally correspond with areas where temperature variability and vulnerability are both relatively high.

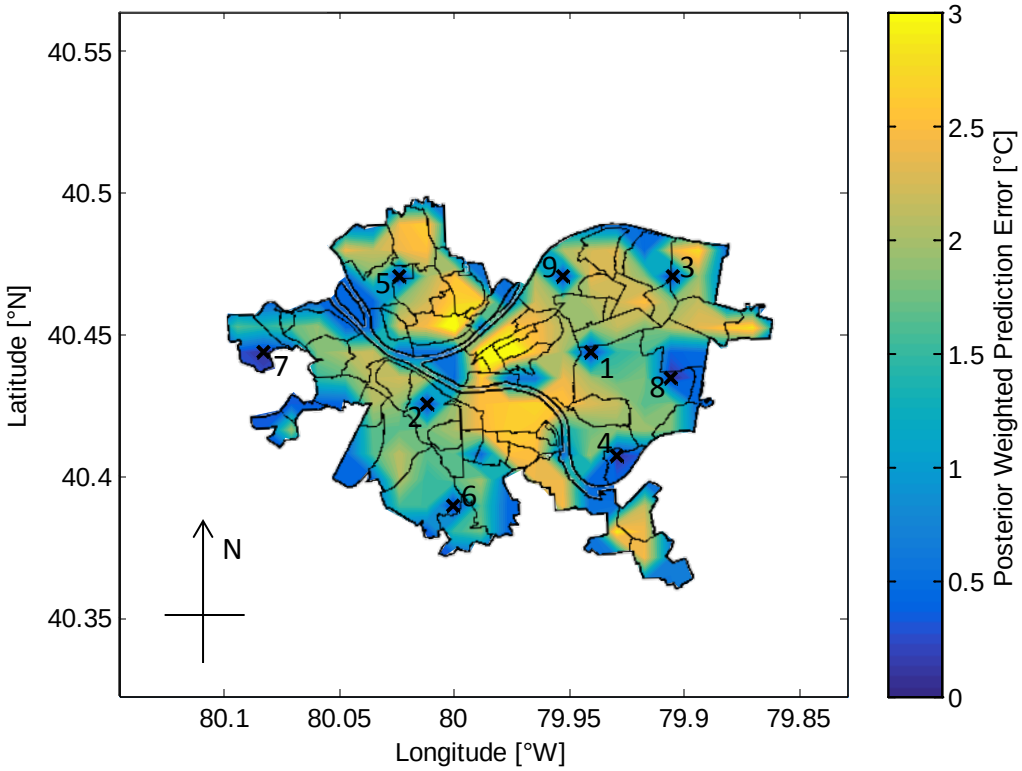

Figure 10-17: Optimized measurement locations for Pittsburgh, PA, supporting the reduction of prediction error of the temperature field weighted by the vulnerability field. Sensing locations are indicated by x's, while the background indicates the posterior weighted prediction error conditional to these measurements. Numbers next to the sensing locations indicate the order the placements were selected by the greedy optimization algorithm. The city's outline is also indicated.

Again, since adding additional sensors will always decrease prediction error, a subset of $Y_{T,All}$ with the same number of sensors as selected under other objectives is used instead. Weighted prediction error gain shows a strong diminishing return property in Figure 10-18, i.e., the slope of the gain decreases as the number of sensors increases. For this reason, a low percentage of the initial value of 6% is chosen to identify the number of sensors to include in the set as 9. Note that this strong diminishing return property contrasts with the weak property exhibited in Figure 10-16. Again, this is due to differences in the definition of the conditional entropy and weighted prediction error metrics, especially the use of the trace instead of the determinant, as well as the logarithmic scaling applied when defining the conditional entropy metric.

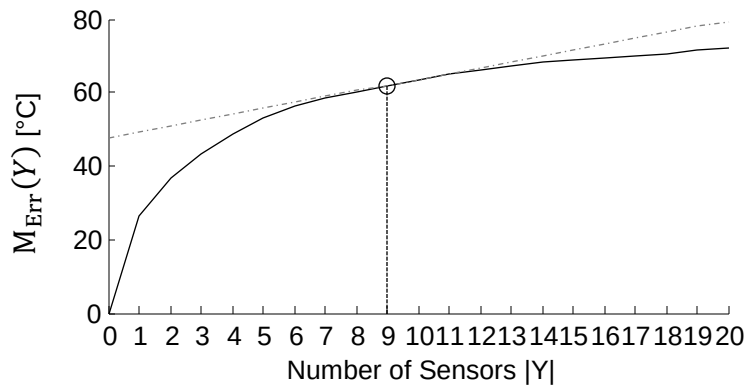

Figure 10-18: Gain in weighted prediction error resulting from the sensor placements through greedy optimization (solid line). The line with slope corresponding to 6% of the gain of the first sensor is also plotted (dashed line), and the point at which this line is tangent to the weighted prediction error gain curve identifies the optimal number of sensors (in this case 9).

10.4.1.3. Value of information for heat wave advisory issuance

Finally, sensor placements are optimized to support the issuance of heat advisories, using the VoI metric based on the decision-making problem outlined in Section 10.3.2.3. The results of this optimization are indicated in Figure 10-19. In this case, a cost for sensing of 0.025 per location (in the same units as $C_f(\mathbf{x})$ and $C_r(\mathbf{x})$) is assumed, representing about 1% of the cost of issuing a heat advisory for that location. Using this cost, the optimal number of sensors is evaluated to be 9. This set is indicated in Figure 10-19, superimposed upon the expected regret (i.e., the difference between the VoI provided by all possible sensors and the value provided by only the selected set of sensors) for a particular date and time, spatially distributed over the city. The VoI and regret are assessed by using the temperature data generated via a numerical weather model for this problem up to this date and time as measurements collected by sensors at the indicated locations and using these data to update the posterior predicted temperature field at this time. Using this updated prediction, an appropriate set of actions for the issuance of heat advisories is chosen. The resulting loss is calculated using the generated temperature field data, and the VoI is calculated for this date and time. The regret is then computed from the VoI. The resulting plot indicates how information collected at just a few locations within the city (less than one tenth of possible measurement locations) allows the predicted temperature field to be updated throughout the area, and for these updated predictions to be used to improve decision-making to reduce losses.

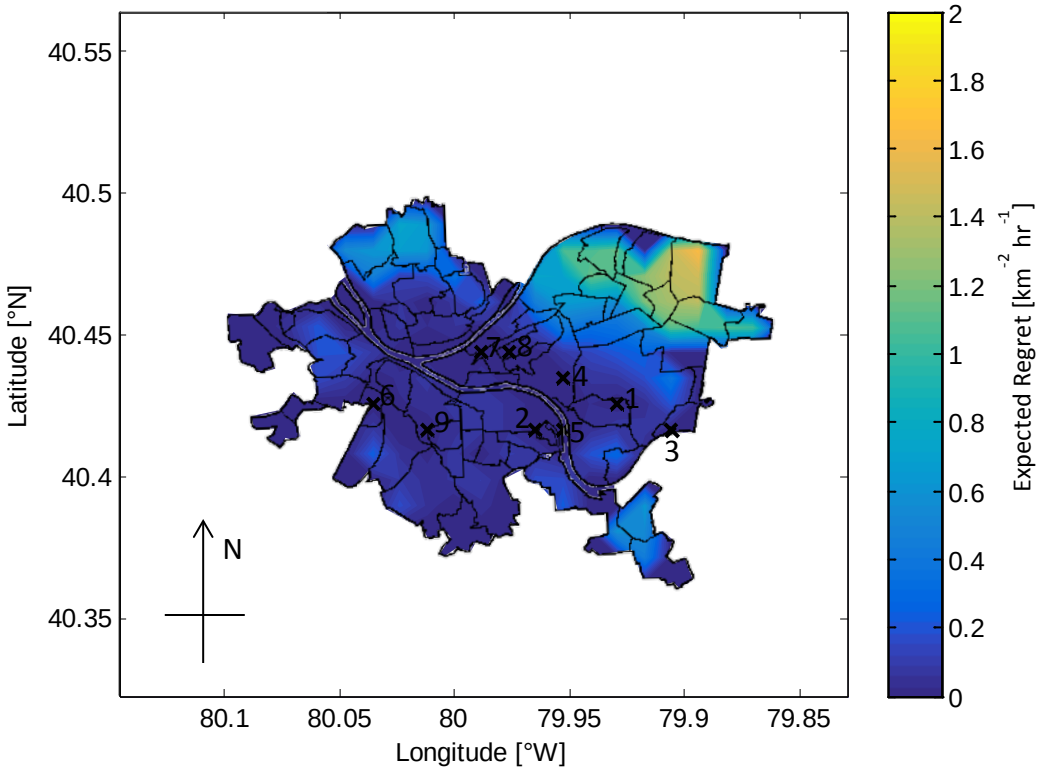

Figure 10-19: Optimized measurement locations for Pittsburgh, PA, supporting decision-making for heat advisory issuance via the VoI metric. Sensing locations are indicated by x's, while the background indicates the expected regret corresponding with the indicated sensor placements, distributed over space, for June 27, 2012 at 1630 hours. Numbers next to the sensing locations indicate the order the placements were selected by the greedy optimization algorithm. The city's outline is also indicated.

For the VoI objective, $\mathrm{VoI}(Y) - \mathrm{C}(Y)$ is plotted directly in Figure 10-20, and the peak location of this curve is chosen as the optimal number of sensors, in this case 9. Note that, for the chosen value of the cost per sensor (0.025), any number of sensors between 2 and 14 would provide roughly the same net VoI; therefore, the optimal number of sensors is quite sensitive to the sensing cost in this case. For example, for a cost per sensor of 0.03, the optimal number of sensors is 4. However, the total VoI remains strictly increasing with the number of sensors (as shown in Figure 10-20 with cost of 0 per sensor).

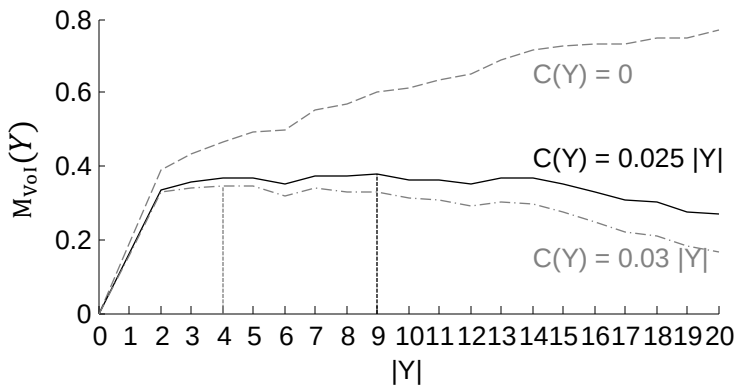

Figure 10-20: VoI resulting from the sensor placements through greedy optimization. Note that the net VoI is maximized when $|Y| = 9$ for a cost of 0.025 per sensor; this is the optimal number of sensors indicated in Figure 10-19. The net VoI under other costs are also shown.

10.4.1.4. Comparison of sensor placements under various metrics

Table 10-H compares the performance of the selected sensor sets under the three objectives considered, i.e., the set $Y^*_{T,Ent}$ optimized under the conditional entropy metric depicted in Figure 10-15, the set $Y^*_{T,Err}$ optimized under the weighted prediction error metric depicted in Figure 10-17, and the set $Y^*_{T,VoI}$ optimized under the VoI metric depicted in Figure 10-19. Performances are listed as percentages of the objective value provided by measuring $Y_{T,All}$; a higher percentage corresponds to a better sensor placement. Note that sets $Y^*_{T,Ent}$, $Y^*_{T,Err}$, and $Y^*_{T,VoI}$ all contain about 7% of the sensors of $Y_{T,All}$. Also note that, as these sensor sets are optimized under the assumption that average temperature predictions Y_{T_0} and cyclic pattern measurements Y_{T_1} are also available, these measurements are included when evaluating the objective. Finally, note that $\mathrm{M_{VoI}}(Y)$ does not include sensing cost in Table 10-H, to allow for easier comparison between the objectives.

Table 10-H: Comparative results for sensor placement metrics.

	Percent of Objective Value for $Y = Y_{T_0} \cup Y_{T_1} \cup Y_{T,All}$		
Sensor Set Y	$\mathrm{M_{Ent}}(Y)$	$\mathrm{M_{Err}}(Y)$	$\mathrm{M_{VoI}}(Y)$
$Y_{T_0} \cup Y_{T_1} \cup Y^*_{T,Ent}$	11.3%	68%	23%
$Y_{T_0} \cup Y_{T_1} \cup Y^*_{T,Err}$	10.8%	72%	35%
$Y_{T_0} \cup Y_{T_1} \cup Y^*_{T,VoI}$	9.7%	57%	42%

As expected, the set optimized under each objective is the best set when evaluated by that objective. Comparing across objectives, the conditional entropy metric does not vary greatly between sets, with the set including $Y^*_{T,VoI}$ performing only slightly worse than the set optimized

under this metric. By contrast, under the VoI metric, the set including $Y^*_{T,Ent}$ achieves just over half the objective value of that including $Y^*_{T,VoI}$. The set including $Y^*_{T,Err}$ always performs either best or second-best, while the set including $Y^*_{T,VoI}$ performs worst except under the VoI metric. This suggests that the weighted prediction error may be a good proxy for the other objectives for performing optimization, but that the VoI metric tends to select sensor placements that may not be those favored by the other objectives.

Comparing sensor placements qualitatively, placements selected under the conditional entropy metric seem to be evenly dispersed across the region, as well as along its perimeter. Placements selected under the weighted predication error metric are similarly distributed, but with fewer measurements along the perimeter. Placements selected under the VoI metric are concentrated in the downtown area, as well as south and southeast of there. Comparing these results with the risk index map of Figure 10-14, there is no clear indication that high-risk or low-risk areas are particularly focused on by any metric. For example, while the VoI metric places some measures in the high-risk downtown area, similarly high-risk areas to the north are ignored, while many sensors are allocated to more moderate-risk areas in the south.

This comparison supports the assertion that heat risk alone is not sufficient to identify areas of interest for temperature monitoring. Some combination of the temperature pattern, prior uncertainty, and vulnerability should be considered. Of the objectives applied here, the VoI takes all of these factors into account through a decision-making problem whose outcomes are defined based on a temperature threshold and whose consequences are defined based on vulnerability. The conditional entropy metric does not take local vulnerability into account, and the weighted prediction error metric does not consider that areas with uncertain but relatively low risk are of less interest for monitoring. Therefore, the VoI metric should be preferred for this problem, as it accounts for relevant decision-making factors that the other metrics do not. However, these other objectives may still be of use in cases where the decision-making problem to be supported is not well-defined or when the goal of the sensing effort is to reduce temperature uncertainty or prediction errors only, without consideration for heat response decision-making implications.

For temperature monitoring in the city of Pittsburgh, this case study suggests two alternative approaches. First, sensing efforts can be distributed across the city; this will best support accurate prediction of temperature throughout, even when factoring in differing vulnerabilities between neighborhoods. Second, monitoring efforts can be focused in central and southern Pittsburgh, as these will best support heat wave advisory issuance under the current decision-making model. These approaches are based on the assumption that information about the city-wide average temperature as well as prior simulations to establish the UHI pattern are available as well; this assumption is relaxed in the following section.

It should be noted, however, that this model may not accurately reflect the priorities of city government and/or the NWS. First, the loss function is modeled based on the vulnerability index developed for the city, which does not necessarily capture the true costs (e.g., monetary impacts and hospitalizations) related to extreme temperature exposure and heat advisory issuance. For

example, the homeless population, which is not generally included in census information and therefore not accounted for by the vulnerability index, will likely be very susceptible to extreme heat, and therefore have a great impact on the heat-related loss function in practice. Second, assumptions about the spatial extent of the area of interest (i.e., the city of Pittsburgh) affect the sensor placement, both in terms of defining the boundaries of the domain and the average temperature considered. In particular, if the region of monitoring interest and the region for which average temperatures are provided by the NWS differ, there will be a systematic difference between these averages that must be corrected for. Third, although it is assumed here that all sensor costs are the same, different costs are possible in practice, both due to the precision of various sensor options and the difficulties of installing sensors in different areas. Greedy placement of sensors with varying costs can also lead to potential suboptimal outcomes via the conditional entropy or weighted prediction error metrics, as discussed in Section 4.4. Finally, although the developed model does seem to adequately represent the temperatures for the city of Pittsburgh, it has only been tested on a set of summer temperatures generated for a single summer, and therefore may not generalize to the temperature distributions for other years. This is especially true if the UHI pattern for the city has undergone a significant change due to major differences in the urban composition. Furthermore, the fact that the most extreme temperatures are not well accounted for by the probabilistic model may influence the performance of sensors under certain atypical summer heat wave scenarios, although this does not appear to be the case given current information about the model performance.

10.4.2. Without Global Mean and Cyclic Pattern Measures

Sensor placement is also performed under the assumption of an information-scarce scenario. The following subsections present results for sensor placement of local temperature measures under the assumption that measures of the global average temperature and the cyclic temperature patterns will not be available (i.e., under use case 1 of Table 10-E). Comparisons are made between these results and those of the previous section, to identify how sensing configurations are affected by this lack of additional information.

10.4.2.1. Conditional entropy of the temperature field

First, sensor placements are optimized based on the conditional entropy metric. The positions of nine greedily optimized measurements are compared in Figure 10-21. Measurements optimized under the assumption that predictions of future average temperatures and data on cyclic temperature patterns would also be available, as presented in Figure 10-15, are denoted as $Y^*_{T,Ent,A}$, while those optimized without the availability of this additional information are denoted as $Y^*_{T,Ent,B}$.

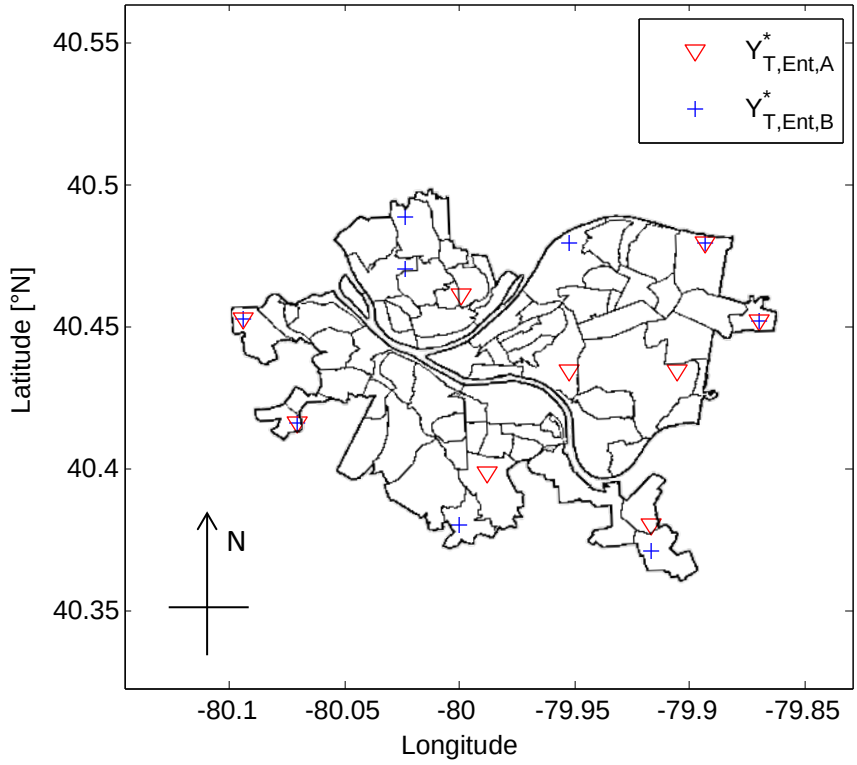

Figure 10-21: Comparison between optimized sensor locations under the conditional entropy metric, with (A) and without (B) global mean and cyclic pattern data available.

Four of nine optimized measurement locations are common between these two sets. For the remaining measurements, those from $Y^*_{T,Ent,B}$ tend to be more widely spaced, i.e., closer to the perimeter of the spatial domain than those from $Y^*_{T,Ent,A}$. This may be a consequence of the fact that correlation distances are slightly shorter for the residual temperature field alone than for the combined residual and cyclic temperature fields.

10.4.2.2. Weighted prediction error of the temperature field

Next, measurements are optimized based on the weighted prediction error metric. A comparison is presented in Figure 10-22 between the set of local measurements optimized under the assumption of additional information as in Figure 10-17, denoted $Y^*_{T,Err,A}$, and that optimized without such additional information, denoted $Y^*_{T,Err,B}$. Here, only one measurement location is common to the two sets. However, qualitatively, the placement patterns are similar, with many measures from one set nearby to their counterparts in the other.

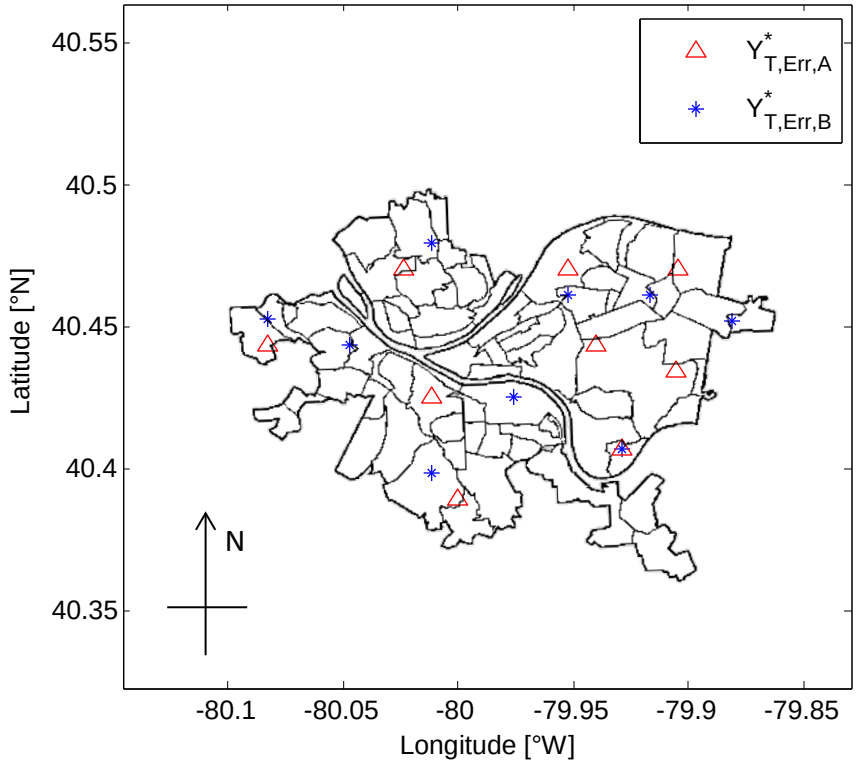

Figure 10-22: Comparison between optimized sensor locations under the weighted prediction error metric, with (A) and without (B) global mean and cyclic pattern data available.

10.4.2.3. Value of information for heat wave advisory issuance

Finally, sensor placements are optimized under the VoI metric, with the sets optimized with ($Y^*_{T,VoI,A}$) and without ($Y^*_{T,VoI,B}$) the assumption of additional information compared in Figure 10-23. In this case, there are no measurements in common between these two sets. Measures in $Y^*_{T,VoI,A}$ tend to be more clustered in the center and south of the city, while those in $Y^*_{T,VoI,B}$ tend to be more evenly distributed, with measures to the northwest and northeast as well. This is probably a consequence of the need to "learn" the cyclic spatial temperature patterns from the data with more widely dispersed sensors in the latter case, whereas data on these patterns are provided as an input when optimizing $Y^*_{T,VoI,A}$.

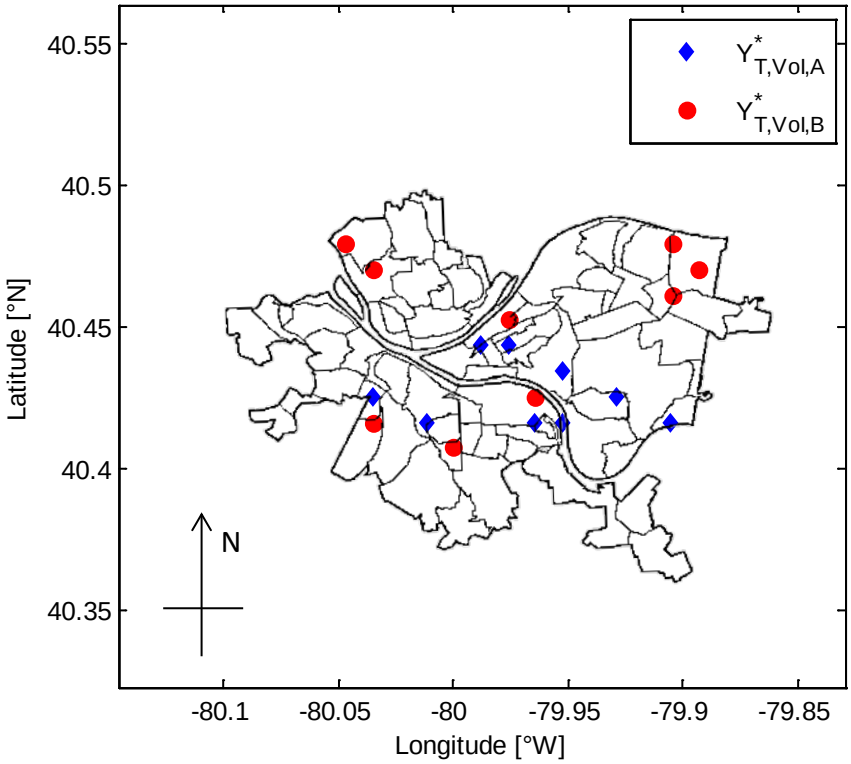

Figure 10-23: Comparison between optimized sensor locations under the VoI metric, with (A) and without (B) global mean and cyclic pattern data available.

For further comparison, the VoI of various sets of measurements are computed and compared in Table 10-I. Note that in Table 10-H VoI values are presented as a percentage of the VoPI considering that additional measurements were also available, i.e., with $\mathrm{M}_{\mathrm{VoI}}\left(Y_{T_0} \cup Y_{T_1} \cup Y_{T,All}\right)$ corresponding to 100% and $\mathrm{M}_{\mathrm{VoI}}\left(Y_{T_0} \cup Y_{T_1}\right)$ corresponding to 0%. In Table 10-I, however, $\mathrm{M}_{\mathrm{VoI}}(\emptyset)$ corresponds to 0%.

Interestingly, an accurate prediction of the regional average temperature alone provides 82% of the VoPI, indicating that even this information is quite valuable by itself in supporting decision-making. This is useful since such regional-averaged temperatures are easily obtainable from regular weather forecasts. A prior assessment of the cyclic temperature pattern provides no value by itself, but when combined with the regional average temperature provides 95% of the VoPI. These cyclic temperature patterns, unlike regional-averaged temperatures, require fine-scale modeling for that particular city, and so are more difficult to obtain.

Table 10-I: Comparative results for VoI of various measurement sets.

Y	$\mathrm{M}_{\mathrm{VoI}}(Y)$	$\mathrm{M}_{\mathrm{VoI}}(Y)/\mathrm{M}_{\mathrm{VoI}}(Y_{T_0} \cup Y_{T_1} \cup Y_{T,All})$
Y_{T_0}	21.31	82%
Y_{T_1}	0.00	0%
$Y_{T_0} \cup Y_{T_1}$	24.67	95%
$Y^*_{T,VoI,A}$	25.02	96%
$Y^*_{T,VoI,B}$	25.03	96%
$Y_{T_0} \cup Y^*_{T,VoI,A}$	25.06	96%
$Y_{T_0} \cup Y^*_{T,VoI,B}$	25.04	96%
$Y_{T_0} \cup Y_{T_1} \cup Y^*_{T,VoI,A}$	25.27	97%
$Y_{T_0} \cup Y_{T_1} \cup Y^*_{T,VoI,B}$	25.13	96%
$Y_{T_0} \cup Y_{T_1} \cup Y_{T,All}$	26.10	100%

As expected, $Y^*_{T,VoI,B}$ slightly outperforms $Y^*_{T,VoI,A}$ in the absence of additional information on the regional average temperature and cyclic temperature pattern. Note that the set of local measurements $Y^*_{T,VoI,A}$ can still be used by itself, even though it is optimized under the assumption that additional information from Y_{T_0} and Y_{T_1} would also be available. Because $Y^*_{T,VoI,A}$ is not being used as intended in this case, its performance is suboptimal; however, the set $Y_{T_0} \cup Y_{T_1} \cup Y^*_{T,VoI,A}$ outperforms $Y_{T_0} \cup Y_{T_1} \cup Y^*_{T,VoI,B}$ when the additional information from Y_{T_0} and Y_{T_1} are available. This is likely because some information gathered as part of $Y^*_{T,VoI,B}$ becomes redundant when measures Y_{T_0} and Y_{T_1} are also available. Note, however, that the VoI provided by both measurement sets is nearly the same; this suggests that the VoI metric is not very sensitive to the placements of local sensors under the assumptions made here in defining the heat advisory decision-making problem for the city of Pittsburgh.

It should also be noted that both $Y^*_{T,VoI,A}$ and $Y^*_{T,VoI,B}$ provide 96% of the VoPI, even without Y_{T_0} and Y_{T_1}. When these additional measures are added, VoI only increases slightly (up to 97% of the VoPI). This shows that, on the one hand, the prior model updated with only local information can account for cyclic temperature patterns and regional temperature trends without direct measurements of these. On the other hand, the direct measurements (i.e., Y_{T_0} and Y_{T_1}) together provide 95% of the VoPI, without any local temperature measures. These results demonstrate that the value provided by local measurements alone is comparable to that provided by an accurate predictive model of the regional temperature combined with a prior fine resolution analysis of the region. This information can be of interest to decision-makers by identifying that these two different temperature modeling activities (i.e., conducting simulations to train local temperature patterns or using sensors to determine these patterns) can have a similar benefit, and therefore only one of these activities need be pursued.

10.4.2.4. Effect of prediction lead times on value of information

Finally, an investigation is made into the effect of prediction lead times on VoI. In the previous sections, VoI is assessed for same-time temperature prediction and warning issuance. That is, decisions about heat advisory issuance are made using information collected up to and including the time of the advisory. If decisions are instead made ahead of time, using a certain lead time, predictions will be less accurate, and therefore provide a lower VoI. Figure 10-24 displays results for how the VoI is affected by the prediction lead time, up to 48 hours. Values are shown for two measurement sets; the optimal set without additional information, $Y^*_{T,VoI,B}$, and the optimal set including additional information, $Y_{T_0} \cup Y_{T_1} \cup Y^*_{T,VoI,A}$. Here, Y_{T_0} represents a forecast for the future average temperature, the standard error of which is assumed to increase linearly with the lead time, such that for 12 hours ahead, the standard deviation of ϵ_{T_0} in Equation 10-17 is 1°C.

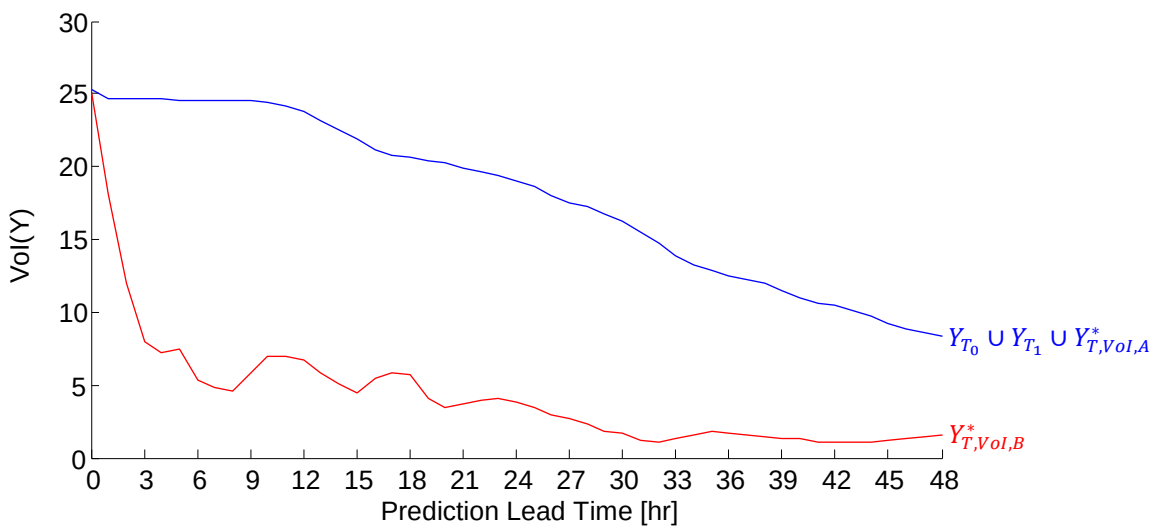

Figure 10-24: VoI versus prediction lead time for two measurement sets.

As expected, VoI tends to decrease with prediction lead time. For measurement set $Y_{T_0} \cup Y_{T_1} \cup Y^*_{T,VoI,A}$, this decrease is gradual, as relatively accurate predictions of the regional average temperature are available to the model, even up to two days ahead. However, without this future forecasting information provided by Y_{T_0}, the VoI for measurement set $Y^*_{T,VoI,B}$ decreases more rapidly with time, dropping below half its initial value at 2 hours ahead, and is about a fifth of its initial value at one day ahead. Interestingly, while there is a high degree of correlation between temperature values at 24 hour intervals, due to cyclic temperature patterns and the daily temperature cycle throughout the region, the VoI does not increase at 24 hours prediction lead time. This is likely because of the temperature variance due to the residual, which exhibits very low correlation at time differences greater than about 12 hours. These results indicate how

important accurate regional temperature predictions can be to ahead-of-time decision-making, as the value provided by local temperature measurements alone decays rather quickly with time.

Chapter 11

Open Issues and Future Work

There remain several key open issues relating to the results presented in this work. First, although throughout this work a greedy optimization approach has been used to minimize the computational issues arising from combinatorial optimization, there remain no guarantees that such an approach will result in optimal or near-optimal sensor placements under objectives such as the VoI. Furthermore, the efficient evaluation of the VoI metric in systems where the loss function is non-decomposable, e.g., series or parallel systems, remains a computational challenge. Finally, issues of sensor placement in problems where the underlying random field structure is unknown or poorly defined have not been addressed here, and alternative approaches to robust optimization and learning in uncertain systems will need to be applied in such problems. However, may opportunities also exist for applying the developed methods for sensor placement in civil engineering, infrastructure system management, and beyond.

11.1. Performance of the Greedy Optimization Approach

As discussed in Section 2.2.1, the greedy optimization approach represents a simple and efficient approximate approach to addressing issues of combinatorial optimization that arise from optimal sensor placement problems. However, this approach does not guarantee an optimal result, and conditions on the near-optimality of results are only provided under certain sensor placement objectives (e.g., conditional entropy). In light of this, the use of such an approach as the basis for the sensor placement results presented throughout this work may appear to be inappropriate and to undermine the validity of the presented results. In addressing this concern, a major open issue (and subject for future work) is assessing the performance of the greedy optimization approach and identifying and implementing alternative approaches to combinatorial optimization that can avoid the shortcomings or pitfalls of the greedy algorithm.

As a first step in this direction, this section will present several numerical examples in which greedily selected sensor placements using the VoI metric are compared to exact placements (i.e., placements optimized through an exhaustive search of the solution space) for tractably small example problems with randomized parameters. In these numerical examples, various parameters defining a system are simulated according to the methods of Table 11-A to define a randomized system. Note that the loss function governing the system is assumed to be decomposable, and is

defined according to the prototypical loss function of Equation 4-15. Further note that the definition of $\mathbf{\Sigma}_F$ and $\mathbf{\Sigma}_\epsilon$ as outlined in the table are designed to ensure the generation of a randomized but symmetrical and positive semi-definite matrix. Simulated systems are considered as having five components, i.e., $n = n_X = 5$. Five potential measurements are also considered.

Table 11-A: Randomized system generation method settings.

Parameter	Equation	Generation Method
$\boldsymbol{\mu}_F$	3-4	$\mu_{f_i} \sim \mathcal{N}(1,1)\ \forall\, i \in \{1,\dots,n_X\}$
$\mathbf{\Sigma}_F$	3-4	$\mathbf{\Sigma}_F = \frac{\mathrm{diag}(\boldsymbol{\sigma}_F)\,\mathbf{K}\,\mathbf{K}^{\mathrm{T}}\,\mathrm{diag}(\boldsymbol{\sigma}_F)}{\mathrm{diag}(\mathbf{K}\,\mathbf{K}^{\mathrm{T}})}$
		where $\log\left(\sigma_{f_i}\right) \sim \mathcal{N}(0,0.5)\ \forall\, i \in \{1,\dots,n_X\}$
		and $\mathbf{K}_{i,j} \sim \mathcal{N}(0,1)\ \forall\, i,j \in \{1,\dots,n_X\}$
$\mathbf{\Omega}_Y$	3-5	$\mathbf{\Omega}_{Y_{i,j}} \sim \mathcal{N}(0,1)\ \ \forall\, i \in \{1,\dots,n_Y\},\ j \in \{1,\dots,n_X\}$
$\boldsymbol{\mu}_\epsilon$	3-6	$\boldsymbol{\mu}_\epsilon = \mathbf{0}$
$\mathbf{\Sigma}_\epsilon$	3-6	$\mathbf{\Sigma}_\epsilon = \frac{\mathrm{diag}(\boldsymbol{\sigma}_\epsilon)\,\mathbf{K}\,\mathbf{K}^{\mathrm{T}}\,\mathrm{diag}(\boldsymbol{\sigma}_\epsilon)}{\mathrm{diag}(\mathbf{K}\,\mathbf{K}^{\mathrm{T}})}$
		where $\log\left(\sigma_{\epsilon_i}\right) \sim \mathcal{N}(-1,0.5)\ \forall\, i \in \{1,\dots,n_Y\}$
		and $\mathbf{K}_{i,j} \sim \mathcal{N}(0,1)\ \forall\, i,j \in \{1,\dots,n_Y\}$
$\mathbf{\Xi}_G$	3-11	$\mathbf{\Xi}_{G_{i,j}} \sim \mathcal{N}(0,1)\ \ \forall\, i \in \{1,\dots,n\},\ j \in \{1,\dots,n_X\}$
$\mathbf{g}_0$	3-11	$\mathbf{g}_0 = \mathbf{0}$
$C_{f,i}$	4-15	$C_{f,i} \sim \mathcal{N}(100,50)\ \forall\, i \in \{1,\dots,n\}$
$C_{r,i}$	4-15	$C_{r,i} \sim \mathcal{N}(50,50)\ \forall\, i \in \{1,\dots,n\}$

For each simulated system, sensor placement is carried out using the VoI metric under both the greedy optimization algorithm and under an exhaustive search algorithm through all 31 possible non-empty sensor placement sets for the systems. The VoI obtained for greedily and exactly optimized sets containing the same numbers of sensors within each system are compared.

The results of this comparison are presented in Figure 11-1, which plots the VoI provided by greedily and exactly optimized sets of the same size. All points lie on or below the identity line, as greedy optimization will always perform as well as or worse than exact optimization. Most points lie on or near the identity line, indicating that in the cases investigated, the greedy algorithm typically performs as well as or nearly as well as an exact, exhaustive approach to sensor placement. In 76% of cases, both algorithms generate sets with the same VoI (although these sets may not necessarily be the same, their performance is comparable). On average, greedily optimized sets provide 98% of the VoI of exhaustively optimized sets of the same size.

Across these 100 simulated systems, in the worst case, the greedily optimized sensor set achieves 45% of the VoI of the exhaustively optimized set.

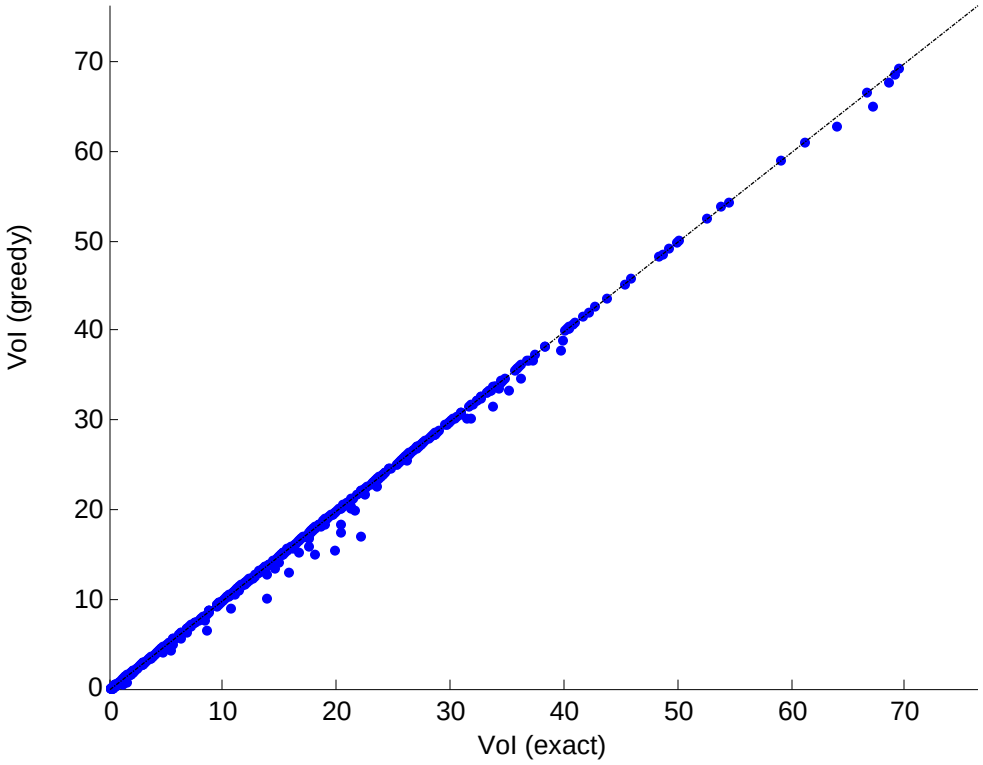

Figure 11-1: Comparative plot of the VoI provided by greedily and exactly optimized measurement sets of the same size across 100 simulated systems. The identity line (dashed) is plotted for reference.

An empirical cumulative distribution function for the ratios between the VoI provided by greedily and exactly optimized sensor sets of the same size across the 100 simulated systems examined is provided in Figure 11-2. This distribution indicates that, across the cases investigated, the greedy algorithm has a high probability of performing nearly as well as an exhaustive search approach (e.g., there is less than a 3% chance of the greedy optimization algorithm generating a set yielding less than 90% of the VoI that might be achieved through an exhaustive search). Finally, note that these results do not signify any general guarantee on the performance of the greedy algorithm, as they only reflect results for systems generated according to the methods outlined in Table 11-A. However, as these methods allow for the generation of a wide variety of different systems (with the notable restriction that the loss function remains of the prototypical decomposable form), it may be reasonable to expect similar results when applying greedy optimization approaches in systems with decomposable loss functions, which covers the majority of cases examined in this work. Therefore, it is expected that the majority of

the results presented here that are obtained through greedy optimization represent reasonable approximations of the true optimal sets for the systems in question.

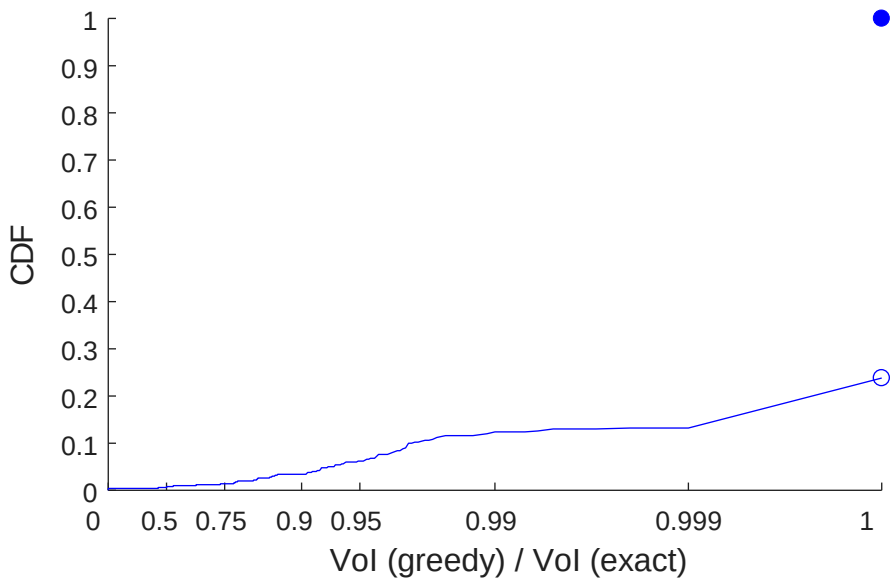

Figure 11-2: Empirical cumulative distribution function for the ratio of the VoI of optimal measurement sets of the same size obtained by greedy and exact optimization. Note the scale of the horizontal axis and discontinuity at 1.

As another investigation into the performance of the greedy optimization algorithm, a comparison of forward and reverse greedy optimization approaches is also undertaken. Recall from Section 2.2.1 that the greedy approach used throughout this work has been the forward greedy approach, in which the optimal sensor set is iteratively built, beginning with an empty set. Conversely, a reverse greedy optimization approach begins with the set $\mathcal{Y}$ of all possible measures and prunes away measures one at a time to arrive at an optimal set. In Section 7.2.3, an example was presented in which the reverse greedy optimization approach represented an improvement over the forward greedy approach in a specific problem case.

Unfortunately, this result will not hold true in general, i.e., the reverse greedy optimization algorithm does not always outperform the forward greedy optimization algorithm. This is verified in a number of numerical simulations, following the same approach outlined above, in which forward and reverse greedy optimization algorithms are compared across a number of randomized systems. The results of this comparison are presented in Figure 11-3.

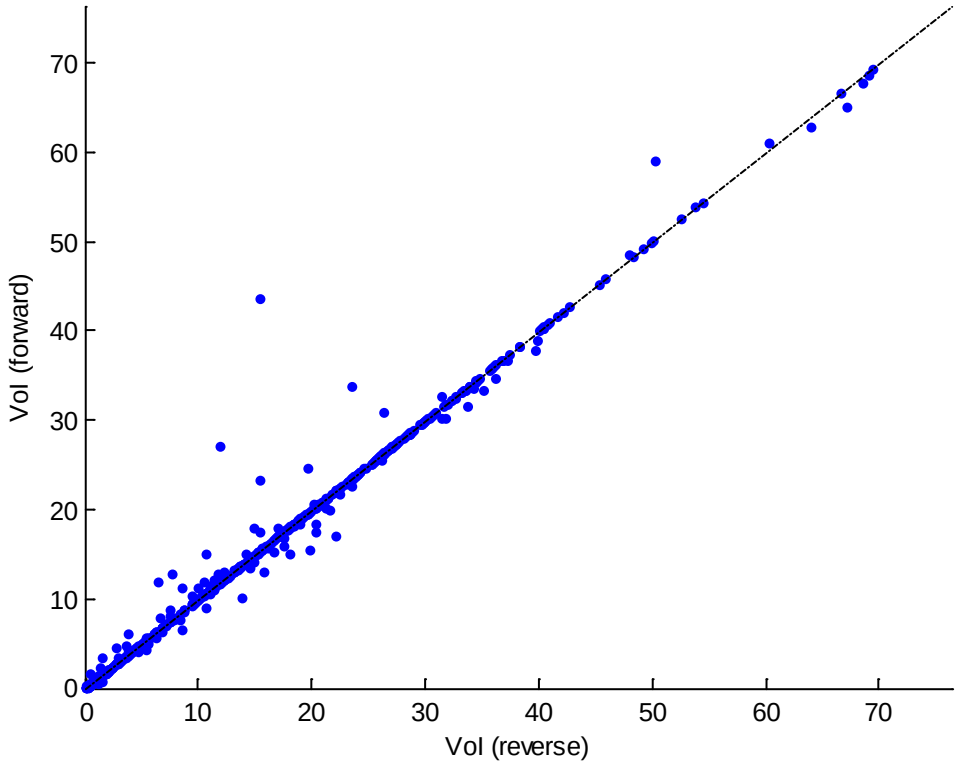

Figure 11-3: Comparative plot of the VoI provided by forward and reverse greedily optimized measurement sets of the same size across 100 simulated systems. The identity line (dashed) is plotted for reference.

As illustrated in Figure 11-3, forward and reverse greedy optimization typically result in selected sensor sets with roughly the same VoI. However, in some cases, either forward or reverse greedy optimization will result in a set with a higher VoI. Across 100 simulated systems investigated, in 24% of cases forward optimization performed better, in 23% of cases reverse optimization performed better, and in the remaining 53% of cases both approaches performed the same. For the two most extreme cases, forward optimization outperformed reverse optimization by 85% (in terms of the difference in VoI of the sets resulting from both approaches, divided by the larger of the two values), and reverse optimization outperformed forward optimization by 55%. However, as can be seen in the figure, for the majority of cases the two approaches do not differ greatly in terms of the VoI of the optimized set, i.e., in 86% of cases one approach outperforms the other by less than 5%. Again, it should be noted that these results are conditional to the random system generation scheme outlined in Table 11-A, and may not apply generally in all situations.

The reason for the use of a forward rather than a reverse greedy optimization approach as the default for this work is due to its computational simplicity. In the majority of cases investigated, the set of all possible measurements is much greater than the set of optimal measurements, i.e.,

$|\mathcal{Y}| \gg |Y^*|$. Therefore, by beginning with the empty set and building towards the optimal set, fewer iterations are required by the forward greedy approach than the reverse.

Various heuristic approaches can also be employed to attempt to avoid suboptimal performance of the forward greedy optimization algorithm in specific applications. One proposed approach is presented in Algorithm 3 below. The motivation for this approach is that, when the rate of increase in the objective value as sensors are greedily selected drops below a certain threshold, the algorithm may find itself in a suboptimal solution path, i.e., it is selecting individual measurements that provide small benefits, while there exist other sets of measurements that when taken together would provide a much larger benefit. Such a case is illustrated in the example of Section 7.2.3, where the forward greedy approach fails to identify pairs of measurements (one measure to correct for bias and one measure to support decision-making) that would provide high value and instead selects individual low-value measures. The proposed algorithm seeks to identify when low-benefit measures are being selected and switch to an alternative or secondary sensor-placement objective that differs from the first and has the potential to "break out" from the suboptimal path being followed by the greedy algorithm. For example, the primary objective might be based on VoI to select informative measures for decision-making, while the secondary objective might be based on mutual information or entropy, in order to select measures that differ greatly from previously selected measures.

Algorithm 3: Pseudo-code for proposed greedy algorithm with alternative heuristic.

Input candidate set $\mathcal{Y}$, primary objective $\mathrm{M}(\cdot)$, secondary objective $\mathrm{M}'(\cdot)$, gain threshold ΔM, cost function $\mathrm{C}(\cdot)$, constraint b
$j = 0,\ \mathcal{Y}_0 = \emptyset$
for each $Y \in \mathcal{Y}$
    if $\mathrm{C}(Y) \leq b$ **then** $\mathcal{Y}_0 = \mathcal{Y}_0 \cup \{Y\}$
end
while $|\mathcal{Y}_j| > 0$
    $j = j + 1$
    select Y_j^* **using** $\mathrm{M}(\cdot)$ (See Equations 2-17 and 2-18)
    if $\mathrm{M}(\{Y_1^*, \dots, Y_j^*\}) - \mathrm{M}(\{Y_1^*, \dots, Y_{j-1}^*\}) < \Delta M$
        select Y_j^* **using** $\mathrm{M}'(\cdot)$
    end
    $\mathcal{Y}_j = \mathcal{Y}_{j-1} \backslash \{Y_j^*\}$
    for each $Y \in \mathcal{Y}_j$
        if $\mathrm{C}(\{Y_1^*, \dots, Y_j^*, Y\}) > b$ **then** $\mathcal{Y}_j = \mathcal{Y}_j \backslash \{Y\}$
    end
end
$j_{\mathrm{end}} = j$
Output $Y^* = \{Y_1^*, \dots, Y_{j_{end}}^*\}$

To demonstrate the application of this proposed approach, a simple problem is examined. Consider two components whose states evolve over ten discrete time steps independently of each other, according to a Gaussian process model with mean function $\mu(\mathbf{x}, t) = 0$ and covariance function:

$$\mathrm{k}(\mathbf{x}_i, t_i, \mathbf{x}_j, t_j) = \sigma(t_i)\,\sigma(t_j)\ \exp\left(-\frac{|t_i - t_j|^2}{2(5)^2}\right)\mathbb{I}(\mathbf{x}_i = \mathbf{x}_j) \tag{11-1}$$

with $\sigma(t) = 0.5 - 0.5\exp(-t)$, modeling independent evolution of these random variables over time domain $\tau = \{0, \ldots, 9\}$. If these state variables exceed a certain threshold, a loss is incurred unless an appropriate corrective action is undertaken. That is, the limit-state variables are defined as in Equation 3-11 with $\mathbf{\Xi}_G = \mathbf{I}$ and $\mathbf{g}_0 = \mathbf{0}$, states are defined as in Equation 7-21 with $g_{\text{upper}} = 0.5$ and $g_{\text{lower}} = -0.5$, and the loss function is spatio-temporally decomposable as in Equation 7-7 with $\gamma(t) = 1$ and with the prototypical local loss functions of Equation 4-15, with $C_f = 100$ and $C_r = 20$. Measurements are possible of the random variable associated with either component at any time step. These measures are very noisy ($\sigma_\epsilon = 100$) but exhibit full correlation of the measurement noise on the same component; this models precise but biased measurements for either component. The measurement cost function is $\mathrm{C}(Y) = 20|Y|$.

For this example, forward greedy sensor selection is performed following Algorithm 3 with the primary VoI-based objective of Equation 4-12 and the secondary entropy-based objective of Equation 4-41. To determine when to employ the secondary objective, in this example $\Delta M = 0$, such that whenever the net additional VoI of the proposed measurement is negative, i.e., the VoI gain from adding the measure is less than the increase in cost, the secondary objective will be used instead to select the next measurement. It should be noted that this next measure will then also lead to a decrease in net VoI, because the optimal next measure by the VoI metric would lead to such a decrease, and therefore any other measure would also lead to the same or a greater decrease. However, in the long-term, the purpose of this heuristic is to select measurements that can later be combined with others to yield a relatively high total net VoI.

Figure 11-4a plots the order in which measures are selected in this example problem via forward and reverse greedy approaches, as well as the proposed heuristic-based approach discussed above. Note that the forward approach, due to the biased measurements, falls into a suboptimal approach of monitoring one component while ignoring the other until all measurements for this first component have been exhausted. The reverse approach largely avoids this, dividing measures evenly between components until eventually abandoning one component to focus on the other, at which point its performance matches that of the forward approach. The heuristic approach follows that of the forward until after the fifth measure is selected. At this point, the gains resulting from adding additional measures fall below the specified threshold, and the next measure is instead chosen by the conditional entropy metric. This leads to a new measure on the second component, which is most different from those selected previously. Moving forward via the greedy approach, new measures are selected for this second component, quickly bringing the

VoI up to the results obtained by the reverse approach, which performs well in this problem. Therefore, it can be seen in this example that the proposed heuristic approach has the potential to avoid suboptimal performance of the forward greedy optimization algorithm under certain conditions, in particular where biased measures must be corrected for. Investigations into the general performance of this proposed approach, including the selection of appropriate secondary objectives under different problem cases and making sure that the secondary objectives properly account for differing measurement costs, remains an area for future work.

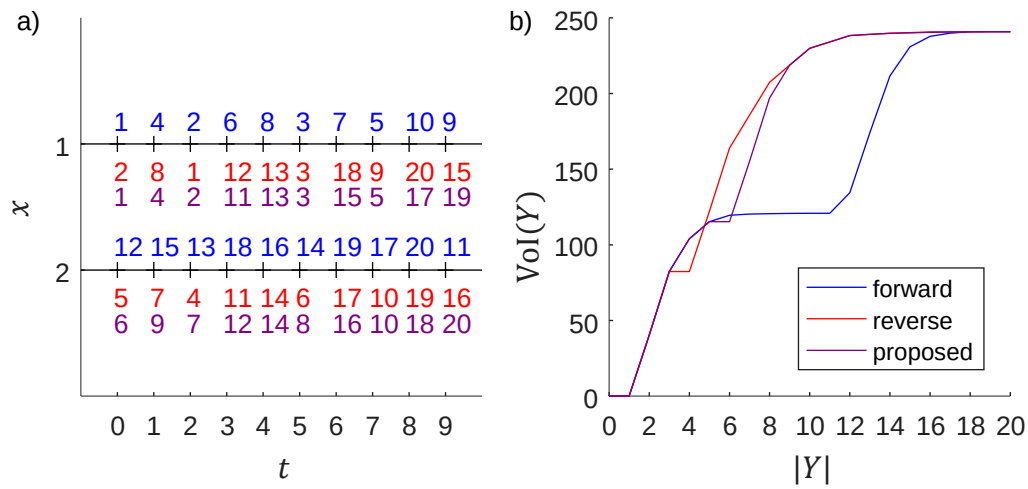

Figure 11-4: Order in which measures are to be included via forward, reverse, and proposed heuristic-based greedy optimization approaches (a) and the VoI of the resulting sets (b).

As noted in Section 1.2.4, there exist a number of alternative approaches to combinatorial optimization, mostly based on heuristics. Any one of these approaches might be applied to optimize the placement of sensors based on the VoI metric. Unfortunately, none of these alternative approaches can provide guarantees on solution optimality either. As observed in Section 7.2.3, expert judgement and knowledge of the underlying problem are often critical to recognizing, understanding, and avoiding potential difficulties with approximate methods such as the greedy optimization approach. It may therefore never be fully possible to remove human decision-making from the optimal sensor placement process, nor should complete automation of this process necessarily be the goal for research. In the end, techniques such as those developed by Mascarenas et al. (2014) whereby computer-based data collection and processing are combined with human sensory and instinctual responses might be the best approach to optimizing sensing for infrastructure system management. For example, a sensor placement optimization "game" might be created (similar to protein bonding games used in biochemistry) in which the VoI metric is used to "score" sensor placements, but human intuition and problem-solving skills are exploited by allowing the player to propose various sensing schemes based on their judgement and knowledge of the system being monitored. The opportunity for effective

human-computer symbiosis to solve complex problems has been identified in many contexts (Sankar, 2012), and the optimization of sensing for infrastructure management may be another example of such a problem.

11.2. Efficient Metric Evaluation in Non-Cumulative Systems

Another major open issue is the efficient evaluation of the VoI metric in large general systems where the cumulative assumption is not valid. As discussed in Section 5.3, in general, evaluation of VoI in large systems is computationally prohibitive, requiring either exhaustive enumeration of system state and management actions or the adoption of approximate techniques that may or may not give sufficiently accurate estimates of VoI. Therefore, the extension of the techniques developed in this work to large systems (with hundreds or thousands of components) in which the system-level loss function is non-decomposable represents a major challenge and obstacle to their implementation. Unfortunately, no simple, generally applicable solutions to this problem were found over the course of this research. Development of such solutions remains an open topic and an area of future research.

In Section 5.4, a heuristic approach to sensor placement in series systems was proposed based on the definition of an equivalent cumulative system. A brief numerical simulation study is conducted to assess the feasibility for this approximate approach. Following the system randomization scheme outlined in Table 11-A, ten systems are simulated. In this case, a series system is defined following the system state definition of Equation 4-18. System failure costs are drawn from a normal distribution with mean 100 and standard deviation 10. After computing the VoI for each potential sensor placement in the series system, the system is converted to an equivalent cumulative system. The loss function of Equation 4-15 is again used, and component failure costs are set equal to the system failure cost for the series system divided by the number of components. All other system parameters are kept the same. VoI is then evaluated for each potential sensor placement in this proposed equivalent cumulative system, and finally scaled by the ratio of the prior expected loss in both systems. Figure 11-5 plots the VoI computed for each sensor set under these two system definitions.

The results indicated in Figure 11-5 can be grouped into three general categories. In the first, indicated in the figure by the blue points, VoI computed for sensor sets in both systems are roughly similar. Deviations from equality can potentially be explained by the relatively low number of simulations used to approximate the VoI in the series system ($n_{sim} = 1000$), such that computations could be completed within a reasonable timeframe. In the second category, indicated by black dots, VoI provided by measurements in the series system is negligible, while VoI in the cumulative system is not. In these cases it is possible that, due to the randomized costs, in the series system the optimal set of actions is always the same regardless of measurement outcomes, while in the cumulative system this is not the case. Finally, in the third category, VoI provided by sensors in both systems seem to follow a roughly linear relationship

(again, accounting for approximation inaccuracies due to the number of samples used). However, this is a different linear relationship from the one predicted in Section 5.4, which was used to set up the series-cumulative system equivalence. In general, these results indicate that, while the approximate method suggested in Section 5.4 can be applicable in some cases, it is not valid in others. One possible reason for this is the relationship between component failure probabilities and management costs. When converting to an equivalent cumulative system, component failure costs are introduced. Relationships between component failure costs, repair costs, and probabilities of failure are thereby introduced which affect the VoI of measurements on each component, as discussed in Section 4.3. Because these relationships did not exist in the series system, introducing them can affect the calculated VoI in the equivalent cumulative system, making it different from that of the series system. Further work will be needed to establish an appropriate equivalence between parameters in series and cumulative systems such that the approximate method suggested in Section 5.4 can be effectively implemented in general.

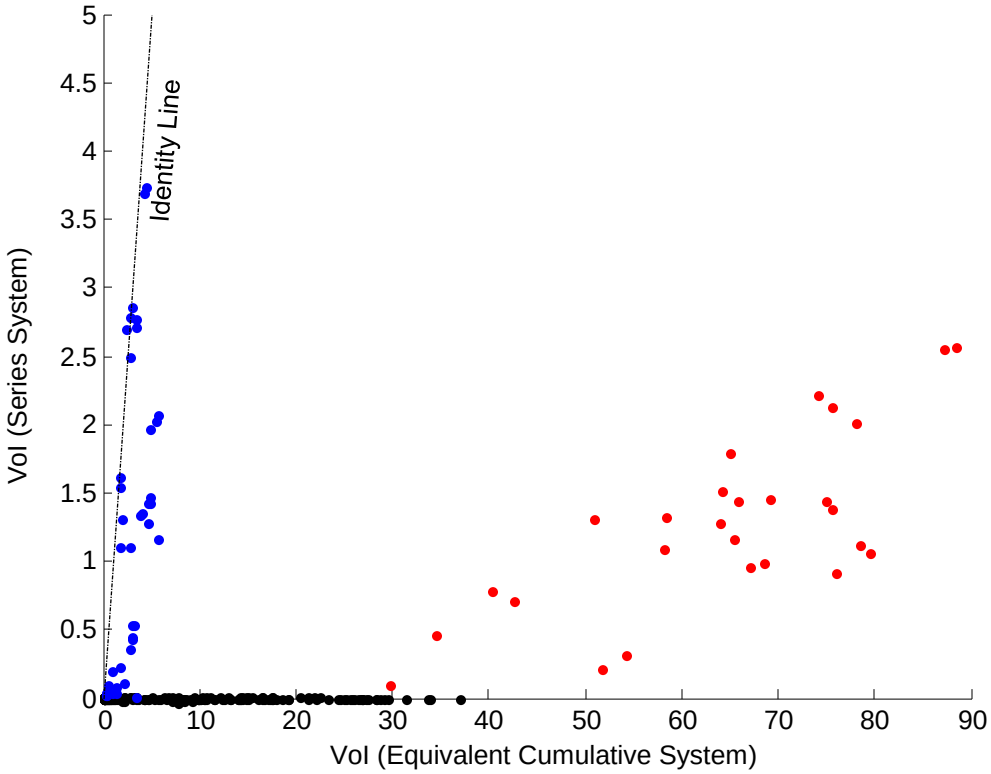

Figure 11-5: Comparative plot of the VoI provided by sets of measurements in series systems and in systems with equivalent cumulative topologies. Results are divided into three groups: sets where the series and equivalent cumulative systems yield similar VoI values (blue), sets where negligible value is provided by sensors in the series system (black), and sets where the VoI obtained in both systems follow an unknown relationship (red). The identity line (dashed) is plotted for reference.

Different system types can also be investigated as well. As an extension of the results presented in Section 5.4, simple series-parallel and parallel-series systems (i.e., systems with multiple series sub-systems in parallel with each other, or with multiple parallel sub-systems in series) can be investigated. It may be possible to investigate these sub-systems using the simplified methods discussed above, and then combine results from different subsystems to evaluate VoI in the overall system. By analyzing more topologically complicated systems in terms of series and parallel sub-systems, it may be possible to approximately analyze a wide variety of different systems using the simplified methods alluded to in this section.

11.3. Uncertainty in Probabilistic Model Definitions

Finally, optimizing inspection plans for systems whose performance (i.e., the static or dynamic relationships that govern their behavior) is uncertain remains an open issue. In Section 2.1.1.1, these properties are grouped together as model parameters W, which might be uncertain, i.e., distributed according to a prior probability distribution. This would represent a hierarchical infrastructure system model, where random field variables F are governed by probabilistic models whose parameters W are themselves random.

A problem of this type is investigated in Chapter 9, where the magnitude and epicenter locations of earthquakes constituted the set of parameters W. Prior distributions for these parameters are assumed (i.e., a homogeneous Poisson process earthquake occurrence model was adopted) and samples are drawn for these parameters. VoI is then evaluated using a system model parameterized by each of these samples and averaging across all samples. It remains to be seen whether this is an appropriate method of accounting for model uncertainty; the resulting sensor placement may not be robust, i.e., it may perform arbitrarily poorly under some parameter settings. Furthermore, for this problem, it is assumed that the parameters, while random, are observable, i.e., that the parameter values are known to the infrastructure managing agent for each realization of an earthquake event. The alternative case, in which these parameters would need to be inferred from collected data, was not investigated.

In general, and especially in situations where sensor placement is performed online, gathered information has the potential to be useful in updating the prior distributions of the model parameters, i.e., in reducing uncertainties in the underlying model as well as in the system itself. Measurement choices must therefore be made that trade off the need to conduct exploration, i.e., learning of the uncertain model parameters, and the need for exploitation, i.e., collecting information to support better decision-making about the system. This problem is explored in Gaussian processes using the conditional entropy and mutual information metric by Krause and Guestrin (2007). Furthermore, the Planning and Learning in Uncertain Systems, or PLUS, approach has been developed by Memarzadeh et al. (2015) and applied to problems of sequential inspection and management in systems whose dynamics are governed by uncertain models. Although developed in the context of Markov deterioration models for systems with discrete states, it is likely that similar approaches can be identified and adapted for use in supporting

optimal sensor placement for the management of infrastructure systems with uncertain model parameters.

Finally, various techniques for robust optimization can be employed in cases of deep uncertainty about the underlying model describing the system. In contrast to approaches that seek to simultaneously learn about the system and support decision-making, these approaches seek only to identify actions or sensor placements that will provide suitable value across a range of possible models for the system. In certain cases, the time or resources to perform simultaneous learning and decision-making may not be available, and sensors must be placed without considering their impact on learning. There are several approaches that might be employed for robust sensing, including placements that maximize the minimum VoI across a range of alternative system models, placements that satisfy some minimum (potentially VoI-based) criteria across all considered models, or placements that allow for the greatest amount of flexibility in sequential sensing and decision-making problems (e.g., Lempert and Collins, 2007). Sensor placements that are robust against uncertainty about the true model for the system will tend to give better performance than placements that are optimized under a particular but incorrect assumption about the underlying probabilistic model.

11.4. Additional Application Cases

Future work should also expand the application case studies considered in order to demonstrate the general applicability of VoI-based sensing optimization. Among the most directly applicable systems for consideration consist of multiple similar components spread out over a spatial domain, for which the overall state of the system can be related to the aggregate states of components. Many systems of interest in civil engineering applications can be categorized in this manner, and several other systems might be simplified to such a system by making certain assumptions.

Several systems relevant to renewable energy, such as wind turbine farms or solar collection centers, represent areas in which the methods developed here might be readily applicable. First, as the overall functionality of the system in terms of generated power can be considered as the sum of the power generated by each individual component, the cumulative system topology assumption is natural for these systems. Second, as these systems consist of a large number of highly similar components distributed over a spatial domain, the use of probabilistic models to perform inference, making use of strong correlations in the performance of similar and/or nearby components, has great potential implications for performing inference from component-level inspection data to system-level knowledge. As these systems become more prevalent, replacing aging energy infrastructure, the need for optimized inspection and maintenance planning for these systems will grow. Furthermore, the benefits of optimized inspection in terms of overall cost reduction can make these newer systems more competitive in an economic sense, thereby accelerating their introduction and the replacement of non-renewable by renewable energy sources.

For the management of structural systems, as discussed in Chapter 8, the main outstanding concern to be addressed in future work is the incorporation of numerous loading conditions and failure modes into a consistent framework for optimizing sensor placement. In accounting for different loading scenarios, e.g., wind loads, snow loads, or earthquake loads, a formulation such as that of Section 9.1.1 might be used, with different distributions over loads defined for each scenario, and sensor placements optimized across all considered scenarios. Alternatively, a common model might be developed to probabilistically characterize structural loadings across all scenarios. Development of such a model would be a complicated endeavor. However, such a model has the potential benefit of assigning correlations between loadings from different scenarios, such that the possibilities of simultaneous extreme loadings due to different scenarios can be accounted for. In terms of system failure modes, the system as a whole can often be modeled as a series arrangement of the failure modes, in that if any one mode occurs, the structure as a whole will fail. This complicates system analysis, preventing a cumulative assumption to be made across the failure modes. However, in structures where there is one predominant mode, or where different portions of the structure operate independently and with different failure criteria, a cumulative assumption might be a reasonable approximation for the analysis of the system.

In terms of managing systems of infrastructures, some systems of interest, including transportation, water supply, and power supply systems, have network structures that strongly influence the performance of the system as a whole. Therefore, typically, the cumulative assumption is not directly applicable, as system-level effects tend to couple decision-making about component management across the system. However, in some cases, reasonable simplifications can be made that allow for approximate analysis of the system under a cumulative topology assumption. For example, in the case of failures resulting from deterioration of components, the event of multiple components simultaneously failing may be extremely unlikely. Therefore, by quantifying failure costs in terms of single component failures only and ignoring system-level effects that can cause the failure of multiple components to have very different costs from the sum of the individual component failure costs, a cumulative assumption might still be reasonably applied to the system. A deeper investigation of such systems and the applicability of this type of assumption is needed to support basing decision-making on it, but intuition suggests there are situations in which such an assumption might lead to nearly correct approximations of measurement VoI.

Finally, in terms of environmental engineering issues such as response to or remediation of contamination or adaptation to changing climate conditions, cumulative assumptions are often reasonable. For example, decisions about contamination are often based on local concentrations, and the overall effectiveness of a remediation effort might be reasonably interpreted as a (possibly weighted) sum of its effectiveness in several discrete areas. In terms of climate adaptation, the fact that local decisions will have little influence on the overall evolution of the climate system allows assumptions such as those presented in Chapter 7 to be applied, thus allowing for more efficient evaluation of the VoI in evolving systems subjected to uncertain

future climate stresses. However, for some applications, the Gaussian process assumption of the underlying variables may be suspect. While in the application of Chapter 10, the Gaussian assumption was reasonable across a wide range of possible temperatures, for other atmospheric factors such as pollutant and particulate matter concentrations, variations of several orders of magnitude and/or multimodal distributions based on different prevailing weather patterns or pollutant source activities may be possible. Such distributions are poorly represented by Gaussian models. More complicated models may need to be used, or non-linear transformations of the variable space may be made to allow for Gaussian process models to describe the transformed variables. More complicated correlation models, e.g. between various atmospheric pollutants based on their chemical interactions, may also be needed.

Chapter 12

Conclusions

This final chapter will recapitulate the major contributions of this work, and summarize the general conclusions that have been drawn.

Beginning in Chapter 1, background material related to infrastructure modeling and monitoring in general was presented, and several major research gaps were identified. First, the lack of a system-level perspective for infrastructure modeling and monitoring can lead to the sub-optimal collection and processing of data on system functioning, in turn leading to sub-optimal decision-making and allocation of resources. Second, the main reason for this lack of a system-level perspective is the computational difficulties associated with scaling approaches to optimal monitoring and management to large systems. Third, there is a general disconnection between management decision-making and monitoring, such that the benefits of information are not typically assessed in terms of how this information will improve decision-making by reducing costs and/or increasing the functionality of the system. In light of these gaps in theory and practice, the primary goal of this work was articulated as the integration of probabilistic modeling and VoI-based optimal sensor placement approaches to efficiently support infrastructure management decision-making.

Towards this goal, Chapter 2 set forth the fundamental structure of the optimal sensor placement problem by defining probabilistic models for infrastructure system behavior, information collection, and management decision-making in general terms. The objective of optimal sensor placement was defined, and the greedy optimization algorithm, used for efficient combinatorial optimization throughout this work, was introduced. Chapter 3 expanded on this basic structure to further define the system behavior, monitoring, and management model in terms of a PGM. Several alternative PGM and influence diagram structures, corresponding to various fundamental assumptions on system behavior and management, were introduced. Notably, the computationally important cumulative system topology was formalized here. Finally, a specific formulation of the problem based on a Gaussian random field model was presented, which represents a flexible and computationally efficient method for system modeling that is used throughout this work.

The VoI metric for sensor placement was introduced in Chapter 4, and the mathematical underpinnings and key properties of this metric were identified. Through a parametric

investigation of the metric in a single-component system, the basic properties of the metric were outlined, namely its dependence on sensor precision, component reliability, and the relative costs of decision-making outcomes. One obstacle to the efficient application of the VoI metric, namely its lack of submodularity, was also identified. Finally, the metric was compared with several alternative metrics for sensor placement, including those based on uncertainty quantification (i.e., conditional entropy) and outcome prediction (i.e., prediction error and misclassification rate). Following this introduction, Chapter 5 focused on the computational difficulties associated with the evaluation of VoI in general, and presented an efficient approach to VoI evaluation in the special case of the cumulative system topology. More generally applicable but approximate methods were also presented for the evaluation of VoI in other systems, together with approximate methods applicable specifically in series and parallel system topologies.

Chapter 6 provided the first complete examples of the application of Gaussian random field models and the VoI metric to support optimal sensor placement. Examples were presented for systems distributed across one- and two-dimensional spatial domains under various assumptions on the system topology. Comparisons were also drawn between the VoI metric and the alternative sensor placement objectives presented in Chapter 4; situations in which these objectives lead to similar or different sensor placement results were identified and discussed.

The basic methodology developed up to this point was extended, in Chapter 7, to spatio-temporal systems whose underlying random features vary in both space and time. Computationally efficient methods for VoI evaluation in such systems, based on the application of the cumulative topology assumption in the temporal domain, were presented. These methods were used to guide optimal sensor placement and scheduling to support long-term infrastructure monitoring and management in several example problems. One situation in particular, in which the structure of the problem led to clearly suboptimal performance by the greedy optimization algorithm, was identified and commented on.

The following three chapters were devoted to presenting practical examples of the application of probabilistic models and the VoI metric to supporting sensor placement and infrastructure system management. In Chapter 8, these methods were applied to a problem of SHM, using the case study of the newly constructed and instrumented Scott Hall building on the CMU campus. In this chapter, information collected by a number of fiber optic strain gauges on the columns of the structure was used to calibrate a probabilistic model of structural behavior and to optimize the selection of a subset of these gauges to support decision-making for a particular potential structural problem. In Chapter 9, a probabilistic model for earthquake occurrence and ground motions in the San Francisco Bay area was used to guide the allocation of SHM activities among multiple components (i.e., bridges and tunnels) in a large transportation infrastructure system to support post-earthquake emergency decision-making. This chapter also demonstrated a case in which the random field parameters of the seismic risk were themselves variable and presented an investigation of the density of VoI across these parameters. Finally, Chapter 10 outlined the development and calibration of a probabilistic model for urban temperatures based on high-

resolution numerical microclimate simulations. This model was then used to optimize the location of temperature sensors to support heat wave advisory issuance in the city of Pittsburgh based on the VoI metric. Various investigations and comparisons of sensor placement performance using different metrics and under various assumptions on the availability of additional information were also presented.

The main body of this work concluded in Chapter 11 by summarizing the major open issues related to optimal sensor placement using the VoI metric, as well as areas for future work in addressing these issues. These issues include verification of the applicability of the greedy optimization approach to the non-submodular VoI metric, where preliminary simulation results indicated that this algorithm may provide near-optimal results across a wide variety of potential systems. However, in certain specific cases, such as when measurements are biased, this algorithm can exhibit very poor performance. An instance of this was illustrated, and a proposed alternative optimization approach, involving an entropy-based heuristic, was outlined. Other open issues included the scaling of VoI evaluation methods in non-cumulative system topologies, as well as the applicability of optimal sensing methods in systems with uncertain parameters. Finally, some possible infrastructures to consider for future work were presented.

Several general, qualitative conclusions can be drawn from the results presented throughout this work. First, applying the various methods presented here, the VoI metric can be used to efficiently optimize the placement of sensors to support management of distributed infrastructure systems, under certain assumptions on the problems structure. Three of these assumptions featured throughout this work are that (i) the performance of the system can be appropriately modeled via a set of jointly Gaussian random variables, (ii) observations taken of the system can be modeled as measurements of linear combinations of these variables, i.e., as variables which themselves share a joint Gaussian distribution, and (iii) the system-level loss function can be expressed as a linear combination of component-level loss functions (with an appropriate definition of the components) such that the loss function is decomposable. The first two assumptions allow for efficient closed-form Bayesian updating of the system model based on observations, and the third allows for efficient computation of the VoI metric itself. As has been demonstrated in this work, there are a wide variety of situations in which these assumptions are valid, and therefore in which the VoI metric can be used efficiently and effectively for sensor placement optimization.

There will, however, always be exceptional cases in which these assumptions are not applicable. For instance, in the probabilistic temperature field model of Chapter 10, Figure 10-10 indicates that the Gaussian model calibrated for temperature does not accurately reflect the true tail distributions for high and low temperatures. Therefore, in this case, the Gaussian random field model, although a reasonable approximation of the true temperature distribution, does not accurately capture the extreme value distribution of the temperature field. In such a case, VoI-based analysis may over- or under-estimate the utility of certain measurements, resulting in a sub-optimal sensor arrangement. Although the impact of this was not considered to be critical for

the applications discussed in Chapter 10, it is nonetheless a concern for the general applicability of the methods presented here. The adoption of appropriate transformations for the underlying Gaussian model, e.g., the use of a log-Gaussian model as in Chapter 9, is one potential approach to overcoming this particular deficiency in practice.

Efficient application of the VoI metric where the loss function is non-decomposable is an additional area for concern, as discussed in Section 11.2. There are several reasons why this assumption might be violated in practice. First, the state of the system may depend in a nonlinear fashion on the states of the components; the classic example, as has been investigated in this work, is that of the series or parallel system topology. Second, actions may have non-local effects; e.g., a single action may affect multiple components or, due to the system topology, an action can affect the overall performance of the system, such as a bridge closure reducing overall traffic capacity in a system. Third, actions may have non-local costs; e.g., there may be economies of scale associated with applying the same action to multiple components, or there may be constraints on how many common actions may be taken such that the cost of performing more than this number of a specific action type throughout a system is prohibitive. All these factors can serve to couple decision-making between system components, and therefore lead the loss function to be non-decomposable. Although there is likely no general and efficient solution for the computational problems presented by non-decomposable loss functions, purely due to the mathematical realities of these problems, various approximate techniques and heuristic methods, such as have been discussed in Sections 5.3 and 11.2, may allow for efficient application of the VoI metric with reasonable expectations of performance in general systems. Additionally, in some cases it may be possible to approximate the true non-decomposable loss function governing a system with an alternative decomposable loss function that can roughly capture the cost structure of the problem in a manner that is more amenable to analysis.

Special mention must also be made of the issues associated with spatio-temporal systems, i.e., systems that change over time as well as are distributed across space. As discussed in Chapter 7, from a theoretical standpoint, analysis of a spatio-temporal system is no different from analysis of a purely spatial system, in that time can be treated as another dimension to the problem. Although this increase of dimensionality leads to a corresponding increase in problem complexity due to the proliferation of system variables and management actions, there is no inherent difficulty associated with the temporal dimension in and of itself. The only caveat to this statement relates to the unidirectional propagation of information in time, i.e., that decisions must be made taking into account only information collected prior to the decision-making time. As with spatial systems, decomposition of the loss function in time can lead to more efficient computability of the VoI metric. In cases where temporal locality does not apply, i.e., where actions can affect the evolution of the system, an alternative efficient approach such as a POMDP framework is needed to avoid the exponential growth in problem complexity inherent with this problem dimensionality increase. Further issues associated with spatio-temporal information gathering, such as the distinction of online and offline sensor placement and the discounting of value for future decision-making, have also been addressed in this work; overall,

accounting for these issues within the general framework presented in this work is relatively straightforward.

From a qualitative perspective, the applicability of the VoI metric to support infrastructure monitoring is clear. Of the various metrics discussed in this work, the VoI alone takes into account the factors of sensor accuracy, prior system knowledge, available management actions, and relative costs or utilities of various outcomes to identify what information will be of the most use in supporting optimal decision-making for system management. Furthermore, the VoI can be directly traded off against sensing cost, as both can be assessed in comparable monetary or utilitarian terms, and so an optimal number of measurements can be selected to maximize the net VoI; for other objectives, a cost-benefit formulation must be adopted to trade off sensing costs with benefits measured in non-monetary terms. Given a well-formulated probabilistic system model and a complete understanding of the loss function governing system management, the VoI metric is, by its definition, the best objective for guiding sensing to support the management of uncertain systems. However, practically, there are a number of obstacles to its widespread usage. The computational barriers discussed previously constitute one obstacle, which may be overcome using the methods discussed in this work wherever applicable. Another obstacle is in the definition of the loss function itself. In some cases, losses may not be clearly understood, or multiple decision objectives may be involved that cannot necessarily be converted easily to a common measure of utility. In these cases, the use of the VoI metric may be suspect due to a lack of fundamental understanding or agreement on the definition of the problem. In such cases, alternative metrics such as the conditional entropy or prediction error metrics may be preferable, as they are agnostic to the decision-making problem structure.

Finally, as discussed in Section 11.1, there are potential shortcomings to the use of the greedy optimization algorithm for any non-submodular metric, including the VoI. In many cases, such as in the example of Section 7.2.3, severe shortcomings of the algorithm can be readily identified though an understanding of the problem structure, e.g., knowing if and where combinations of measurements can have emergent properties that make them more useful than simple aggregation would suggest. Throughout this work, the greedy approach has been used for the optimization of sensor placements based on the VoI metric, and wherever the optimality of these placements are assessed (either through comparison with an exactly optimized sensor placement, with randomly selected sensors, or with the VoPI) the greedily optimized placement has represented a reasonable approximation to the true optimal or best possible solution. Based on this empirical evidence, this work posits that the use of the greedy optimization algorithm can be justified to support optimal sensor placement for infrastructure system management using the VoI metric. Further empirical investigations following the preliminary results of Section 11.1 or theoretical results as presented in Appendix B may provide additional support for this claim in the future.

In summary, sensor placement for infrastructure management is best supported by the VoI metric, which implicitly accounts for problem uncertainty and potential outcomes to determine the utility of measurements in supporting decision-making. This metric can be efficiently

evaluated in large systems under certain assumptions on the problem structure. Where these assumptions do not apply, approximate methods may be utilized. This work presents many examples of the application of the VoI metric, together with probabilistic models of uncertain systems and a forward greedy measurement selection algorithm, in which optimized sets of sensor locations and/or inspections can be efficiently selected to support the management of a distributed system. Other metrics have also been investigated for comparative purposes, and to serve as alternatives in cases where the underlying problem structure is too ill-defined to allow for computation of the VoI metric. Overall, this work has presented a simple, scalable, efficient method for optimal sensor placement to support infrastructure system management based on system-level PGMs and the VoI metric.

12.1. Related Publications

This section lists the publications related to this work, and what major contributions are discussed in each publication.

Preliminary investigations evaluating the conditional entropy metric in multi-component systems are presented by Malings et al. (2013). This paper identifies some general issues relating to the use of the conditional entropy metric in supporting optimal monitoring of infrastructure systems. Although not directly relevant to this work, these results help to motivate the discussion of Section 1.2.5 regarding the shortcomings of other sensing metrics.

The formulation of the VoI metric using a decomposable loss function, as in Equation 3-1, was first presented by Malings and Pozzi (2014a). Also first published in that paper were the application of the VoI metric to the problem of seismic risk monitoring, as discussed in Chapter 9, the efficient computational techniques of Section 5.2.2, and a preliminary formulation of the seismic risk model of Section 9.1. These preliminary results were expanded by Malings and Pozzi (2014b) to consider continuous fault lines, as in Section 9.2 (whereas previously a set of discrete epicenter locations was considered) and to present an analysis of the VoI density for the seismic risk problem. A further extension and summarization of this work was presented by Malings and Pozzi (2016a); this provided a complete summary of the previously published work on seismic risk analysis and optimal sensing using VoI, as presented in Chapter 9, including a previously omitted formulation for the evaluation of VoI for a Poisson process and under earthquake scenario uncertainty. Furthermore, this paper presented a qualitative comparison between the VoI and conditional entropy metrics, which forms a portion of Section 4.4. Finally, a parametric analysis of the VoI in multicomponent cumulative systems, with an application to seismic risk, was presented by Malings and Pozzi (2015). Although the results of this analysis are not directly included in this work, these did help to motivate some of the discussions included in Chapter 9.

An analysis of the complexity of VoI computation, as discussed in Section 5.1, was first presented by Malings and Pozzi (2016b). Other significant contributions of that work included comparisons between the VoI, prediction, and misclassification metrics, as presented in Section

4.4, a parametric analysis of the VoI for a single component, as presented in Section 4.3, efficient computational techniques for special cases of cumulative system topologies, as presented in Section 5.2.1, general approximate computational techniques for VoI evaluation, as presented in Section 5.3, and comparative analyses between various metrics in one- and two-dimensional problems, as presented in Chapter 6. An analysis of the performance of the VoI metric in multicomponent non-cumulative series and parallel systems was first presented by (Malings and Pozzi, 2016c), along with the example presented in Section 5.4.

The extension of the VoI metric and optimal sensor placement approaches to spatio-temporal systems as discussed in Chapter 7 was first presented by Malings and Pozzi (2016d). This publication included the formulation of VoI in spatio-temporal systems as in Section 7.1, the use of temporally decomposable loss functions to efficiently evaluate VoI for these systems (analogous to the use of decomposable loss functions in purely spatial problems, as discussed in previous works) and the demonstrative examples presented in Sections 7.2.1 and 7.2.3 (the example of Section 7.2.2 was taken from an earlier unpublished version of this paper).

The fiber optic sensing system installed for the monitoring of the Scott Hall building, as discussed in Chapter 8, as well as the optimal sensor placement example based on data collected by these sensors, as in Section 8.3, were first discussed by Malings et al. (2015). Follow-on results from this project, particularly the optimal sensor placement and scheduling results presented in Section 8.4, were presented by Malings and Pozzi (2016e).

The results of the SHADE (surface heat assessment for developed environments) project, as summarized in Chapter 10, were presented in several publications. The probabilistic urban temperature model outlined and validated in Section 10.2 was introduced by Malings et al. (2017a). This model was used for optimal allocation of temperature monitoring efforts, as discussed in Sections 10.3 and 10.4, by Malings et al. (2017b). A further analysis of optimal sensing for urban temperatures, under various alternative assumptions on the prediction lead-time and availability of information, as in Section 10.4.2, was presented by Malings et al. (2016).

References

Anagnos, T., Kiremidjian, A., 1988. A review of earthquake occurrence models for seismic hazard analysis. Probabilistic Engineering Mechanics 3, 3–11.

Anderson, B.D.O., Moore, J.B., 2007. Optimal control: linear quadratic methods, Dover ed., augmented republ. of the ed. 1990. ed, Dover books on engineering. Dover Publications, Mineola, NY.

Anderson, G.B., Bell, M.L., 2011. Heat Waves in the United States: Mortality Risk during Heat Waves and Effect Modification by Heat Wave Characteristics in 43 U.S. Communities. Environmental Health Perspectives 119, 210–218. doi:10.1289/ehp.1002313

Andrews, K., 1994. The consequences of heatwaves in Australia (Hons. thesis). School of Earth Sciences, Macquarie University, Sydney, Australia.

Ariaratnam, S.T., El-Assaly, A., Yang, Y., 2001. Assessment of Infrastructure Inspection Needs Using Logistic Models. Journal of Infrastructure Systems 7, 160–165. doi:10.1061/(ASCE)1076-0342(2001)7:4(160)

ASCE, 2016. Failure to Act: closing the infrastructure investment gap for America's economic future. American Society of Civil Engineers; Economic Development Research Group, Inc.

ASCE, 2013. Report Card for America's Infrastructure.

Au, S.-K., Beck, J.L., 2001. Estimation of small failure probabilities in high dimensions by subset simulation. Probabilistic Engineering Mechanics 16, 263–277. doi:10.1016/S0266-8920(01)00019-4

Aubrecht, C., Özceylan, D., 2013. Identification of heat risk patterns in the U.S. National Capital Region by integrating heat stress and related vulnerability. Environment International 56, 65–77. doi:10.1016/j.envint.2013.03.005

Badea, G.E., Cret, P., Lolea, M., Setel, A., 2011. Studies of Carbon Steel Corrosion in Atmospheric Conditions. ACTA Technica Corviniensis - Bulletin of Engineering 4.

Baker, J.W., Miller, M.K., Markhvida, M., 2015. Local Measures of Disruption for Quantifying Seismic Risk and Reliability of Complex Networks, in: Proceedings of the 12th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP12). Vancouver, Canada.

Bednarz, E.T., Zhu, W.D., Smith, S.A., 2013a. Identifying Magnitudes and Locations of Multiple Loads and the Resultant of a Distributed Load on a Slender Beam Using Strain Gage Based Methods: Load Identification Using Strain Gage Based Methods. Experimental Techniques. doi:10.1111/ext.12036

Bednarz, E.T., Zhu, W.D., Smith, S.A., 2013b. Identifying the Magnitude and Location of a Load on a Slender Beam Using a Strain Gauge Based Force Transducer: Load Identification Using a Strain Gauge Based Force Transducer. Strain 49, 276–285. doi:10.1111/str.12035

Bellman, R., 1984. Dynamic programming. Princeton Univ. Pr, Princeton, NJ.

Bensi, M., Der Kiureghian, A., Straub, D., 2011. A Bayesian Network Methodology for Infrastructure Seismic Risk Assessment and Decision Support (PEER No. PEER Report 2011/02). Pacific Earthquake Engineering Research Center, College of Engineering, University of California, Berkeley, Berkeley, California, USA.

Bensi, M., Kiureghian, A.D., Straub, D., 2015. Framework for Post-Earthquake Risk Assessment and Decision Making for Infrastructure Systems. ASCE-ASME Journal of Risk and

Uncertainty in Engineering Systems, Part A: Civil Engineering 1, 04014003. doi:10.1061/AJRUA6.0000810

Bensi, M., Kiureghian, A.D., Straub, D., 2013. Efficient Bayesian network modeling of systems. Reliability Engineering & System Safety 112, 200–213. doi:10.1016/j.ress.2012.11.017

Berger, J.O., 1993. Statistical decision theory and Bayesian analysis, 2nd ed, Springer series in statistics. Springer-Verlag, New York.

Berko, J., Ingram, D.D., Saha, S., Parker, J.D., 2014. Deaths attributed to heat, cold, and other weather events in the United States, 2006– 2010 (No. 76), National health statistics reports. National Center for Health Statistics, Hyattsville, MD.

Bertelè, U., Brioschi, F., 1972. Nonserial Dynamic Programming. Academic Press, New York, NY.

Borgatti, S.P., 2006. Identifying sets of key players in a social network. Computational and Mathematical Organization Theory 12, 21–34. doi:10.1007/s10588-006-7084-x

Bosello, F., Roson, R., Tol, R.S.J., 2007. Economy-wide Estimates of the Implications of Climate Change: Sea Level Rise. Environmental and Resource Economics 37, 549–571. doi:10.1007/s10640-006-9048-5

Boyd, S.P., Vandenberghe, L., 2004. Convex optimization. Cambridge University Press, Cambridge, UK.

Bradford, K., Abrahams, L., Hegglin, M., Klima, K., 2015. A Heat Vulnerability Index and Adaptation Solutions for Pittsburgh, Pennsylvania. Environ. Sci. Technol. 49, 11303–11311. doi:10.1021/acs.est.5b03127

Brandimarte, P., 2006. Non-convex optimization, in: Numerical Methods in Finance and Economics. John Wiley & Sons, Inc, Hoboken, NJ, USA, pp. 563–599.

Brunner, E., 2001. Commentary: Education, education, education. International Journal of Epidemiology 30, 1126–1128. doi:10.1093/ije/30.5.1126

Buldyrev, S.V., Parshani, R., Paul, G., Stanley, H.E., Havlin, S., 2010. Catastrophic cascade of failures in interdependent networks. Nature 464, 1025–1028. doi:10.1038/nature08932

Buscail, C., Upegui, E., Viel, J.-F., 2012. Mapping heatwave health risk at the community level for public health action. International Journal of Health Geographics 11, 38. doi:10.1186/1476-072X-11-38

Büyüköztürk, O., NDTMS, International Union of Laboratories and Experts in Construction Materials, Systems and Structures (Eds.), 2013. Nondestructive testing of materials and structures: proceedings of NDTMS-2011, Istanbul, Turkey, May 15 - 18, 2011 ; [International Symposium on Nondestructive Testing of Materials and Structures], RILEM bookseries. Springer, Dordrecht.

Byon, E., Ding, Y., 2010. Season-Dependent Condition-Based Maintenance for a Wind Turbine Using a Partially Observed Markov Decision Process. IEEE Transactions on Power Systems 25, 1823–1834. doi:10.1109/TPWRS.2010.2043269

Cadini, F., Zio, E., Petrescu, C.-A., 2009. Using Centrality Measures to Rank the Importance of the Components of a Complex Network Infrastructure, in: Setola, R., Geretshuber, S. (Eds.), Critical Information Infrastructure Security. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 155–167.

Casella, G., Berger, R.L., 2002. Statistical inference, 2nd ed. ed. Thomson Learning, Australia ; Pacific Grove, CA.

Chen, F., Zhang, Y., 2009. On the coupling strength between the land surface and the atmosphere: From viewpoint of surface exchange coefficients. Geophysical Research Letters 36. doi:10.1029/2009GL037980

Chen, Y., Bagnell, J.A., Javdani, S., Srinivasa, S., Karbasi, A., Krause, A., 2015. Submodular surrogates for value of information, in: Proceedings of the 29th AAAI Conference on Artificial Intelligence.

Chen, Y., Krause, A., 2013. Near-optimal Batch Mode Active Learning and Adaptive Submodular Optimization, in: Proceedings of the 30th International Conference on Machine Learning. Atlanta, Georgia, USA.

Clark, A., Alomair, B., Bushnell, L., Poovendran, R., 2016. Submodularity in Dynamics and Control of Networked Systems, Communications and Control Engineering. Springer International Publishing, Cham.

Claxton, K., Posnett, J., 1996. An economic approach to clinical trial design and research priority-setting. Health Economics 5, 513–524. doi:10.1002/(SICI)1099-1050(199611)5:6<513::AID-HEC237>3.0.CO;2-9

Clougherty, J.E., Kubzansky, L.D., 2009. A Framework for Examining Social Stress and Susceptibility to Air Pollution in Respiratory Health. Environmental Health Perspectives 117, 1351–1358. doi:10.1289/ehp.0900612

Colorini, A., Dorigo, M., Maffioli, F., Maniezzo, V., Righini, G., Trubian, M., 1996. Heuristics from Nature for Hard Combinatorial Optimization Problems. International Transactions in Operational Research 3, 1–21.

Cook, W.J., Cunningham, W.H., Pulleyblank, W.R., Schrijver, A. (Eds.), 1998. Combinatorial optimization, Wiley-Interscience series in discrete mathematics and optimization. Wiley, New York.

Cover, T., Thomas, J., 2006. Elements of Information Theory, 2nd ed. John Wiley & Sons, Inc., Hoboken, New Jersey, USA.

Cressie, N.A.C., Wikle, C.K., 2011. Statistics for spatio-temporal data, Wiley series in probability and statistics. Wiley, Hoboken, N.J.

Crichton, D., 1999. The Risk Triangle, in: Ingleton, J. (Ed.), Natural Disaster Management. Tudor Rose, London, UK, pp. 102–103.

Curriero, F.C., Heiner, K.S., Samet, J.M., Zeger, S.L., Strug, L., Patz, J.A., 2002. Temperature and Mortality in 11 Cities of the Eastern United States. American Journal of Epidemiology 155, 80–87. doi:10.1093/aje/155.1.80

Currin, C., Mitchell, T., Morris, M., Ylvisaker, D., 1991. Bayesian prediction of deterministic functions, with applications to the design and analysis of computer experiments. Journal of the American Statistical Association 86, 953–963.

Dam Safety in the United States: A Progress Report on the National Dam Safety Program, 2009. . Federal Emergency Management Agency.

Das, A., Kempe, D., 2008. Algorithms for subset selection in linear regression. ACM Press, p. 45. doi:10.1145/1374376.1374384

Das, B.M., 2010. Principles of geotechnical engineering. Cengage Learning, Stamford, Conn.

Davis, R., Knappenberg, P., Michaels, P., Novicoff, W., 2003. Changing Heat-Related Mortality in the United States. Environmental Health Perspectives 111, 1712–1718.

Dehghani, M.S., Flintsch, G., McNeil, S., 2014. Impact of Road Conditions and Disruption Uncertainties on Network Vulnerability. Journal of Infrastructure Systems 20, 04014015. doi:10.1061/(ASCE)IS.1943-555X.0000205

Ditlevsen, O., Madsen, H.O., 1996. Structural Reliability Methods, 1st ed. John Wiley & Sons Ltd, Chichester, UK.

Douglas, J., 2011. Ground-motion prediction equations, 1964-2010 (PEER No. PEER Report 2011/102). Pacific Earthquake Engineering Research Center, College of Engineering, University of California, Berkeley, Berkeley, California, USA.

Doyon, B., Bélanger, D., Gosselin, P., 2008. The potential impact of climate change on annual and seasonal mortality for three cities in Québec, Canada. International Journal of Health Geographics 7, 23. doi:10.1186/1476-072X-7-23

Dudhia, J., 1989. Numerical Study of Convection Observed during the Winter Monsoon Experiment Using a Mesoscale Two-Dimensional Model. Journal of the Atmospheric Sciences 46, 3077–3107. doi:10.1175/1520-0469(1989)046<3077:NSOCOD>2.0.CO;2

Dueñas-Osorio, L., Craig, J.I., Goodno, B.J., Bostrom, A., 2007. Interdependent Response of Networked Systems. Journal of Infrastructure Systems 13, 185–194. doi:10.1061/(ASCE)1076-0342(2007)13:3(185)

Endrenyi, J., Aboresheid, S., Allan, R.N., Anders, G.J., Asgarpoor, S., Billinton, R., Chowdhury, N., Dialynas, E.N., Fipper, M., Fletcher, R.H., Grigg, C., McCalley, J., Meliopoulos, S., Mielnik, T.C., Nitu, P., Rau, N., Reppen, N.D., Salvaderi, L., Schneider, A., Singh, C., 2001. The present status of maintenance strategies and the impact of maintenance on reliability. IEEE Transactions on Power Systems 16, 638–646. doi:10.1109/59.962408

Faber, M.H., Koehler, J., Nishijima, K. (Eds.), 2011. Applications of statistics and probability in civil engineering: proceedings of the 11th International Conference on Applications of Statistics and Probability in Civil Engineering, Zürich, Switzerland, 1 - 4 August 2011, A Balkema book. CRC Press, Boca Raton, Fla.

Fan, W., Qiao, P., 2011. Vibration-based Damage Identification Methods: A Review and Comparative Study. Structural Health Monitoring 10, 83–111. doi:10.1177/1475921710365419

Farrar, C.R., 2013. Structural health monitoring: a machine learning perspective. Wiley, Chichester, West Sussex, U.K. ; Hoboken, N.J.

Field, E.H., Dawson, T.E., Felzer, K.R., Frankel, A.D., Gupta, V., Jordan, T.H., Parsons, T., Petersen, M.D., Stein, R.S., Weldon II, R.J., Wills, C.J., 2009. Uniform California Earthquake Rupture Forecast, Version 2 (UCERF 2). Bulletin of the Seismological Society of America 99, 2053–2107. doi:10.1785/0120080049

Flynn, E.B., Todd, M.D., 2010. A Bayesian approach to optimal sensor placement for structural health monitoring with application to active sensing. Mechanical Systems and Signal Processing 24, 891–903. doi:10.1016/j.ymssp.2009.09.003

Galvan, G., Agarwal, J., 2015. Vulnerability analysis of interdependent infrastructure systems. doi:10.14288/1.0076297

Garré, L., Friis-Hansen, P., 2013. Using Bayesian Networks and Value of Information for riskbased adaptation to climate change: an application of the DNV-ADAPT framework, in: Proceedings of ICOSSAR 2013. New York, NY.

Geerts, Linacre, 1999. Data from a newsletter of the Natural Hazards Research Centre. Macquarie University, Sydney, Australia.

Gertsbakh, I.B., Shpungin, Y., 2011. Network reliability and resilience. Springer, Heidelberg; New York.

Glisic, B., 2015. The Sixth Short Course on Structural Health Monitoring using Fiber Optic Sensors.

Glisic, B., Inaudi, D., Lau, J., Mok, Y., Ng, C., 2005. Long-term monitoring of high-rise buildings using long-gage fiber optic sensors, in: Proceedings of 7th International Conference on Multi-Purpose High-Rise Towers and Tall Buildings. Dubai, UAE.

Golovin, D., Krause, A., 2011. Adaptive Submodularity: Theory and Applications in Active Learning and Stochastic Optimization. Journal of Artificial Intelligence Research 42, 427–486.

González, A.D., Dueñas-Osorio, L., Sánchez-Silva, M., Medaglia, A.L., 2015. The computational complexity of probabilistic Interdependent Network Design Problems. doi:10.14288/1.0076118

Gotovos, A., Casati, N., Hitz, G., Krause, A., 2013. Active Learning for Level Set Estimation, in: Proceedings of the 23rd International Joint Conference on Artificial Intelligence. Presented at the 23rd International Joint Conference on Artificial Intelligence, Beijing, China.

Goulet, J.-A., Kiureghian, A.D., Li, B., 2015a. Pre-posterior optimization of sequence of measurement and intervention actions under structural reliability constraint, in: Proceedings of the 12th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP12). Vancouver, Canada.

Goulet, J.-A., Kiureghian, A.D., Li, B., 2015b. Pre-posterior optimization of sequence of measurement and intervention actions under structural reliability constraint. Structural Safety 52, 1–9. doi:10.1016/j.strusafe.2014.08.001

Grower, S., Mee, C., Campbell, M., 2011. Protecting Vulnerable People from Health Impacts of Extreme Heat. Toronto Public Health, Toronto, Ontario, Canada.

Guestrin, C., Krause, A., Singh, A., 2005. Near-Optimal Sensor Placements in Gaussian Processes, in: Proceedings of the 22nd International Conference on Machine Learning. Presented at the 22nd International Conference on Machine Learning, Bonn, Germany.

Guo, H.Y., Zhang, L., Zhang, L.L., Zhou, J.X., 2004. Optimal placement of sensors for structural health monitoring using improved genetic algorithms. Smart Materials and Structures 13, 528–534. doi:10.1088/0964-1726/13/3/011

Gutenberg, B., Richter, C.F., 1954. Seismicity of the Earth and Associated Phenomena, 2nd ed. Princeton University Press, Princeton, NJ.

Hajat, S., Kovats, R.S., Lachowycz, K., 2006. Heat-related and cold-related deaths in England and Wales: who is at risk? Occupational and Environmental Medicine 64, 93–100. doi:10.1136/oem.2006.029017

Harlan, S.L., Declet-Barreto, J.H., Stefanov, W.L., Petitti, D.B., 2012. Neighborhood Effects on Heat Deaths: Social and Environmental Predictors of Vulnerability in Maricopa County, Arizona. Environmental Health Perspectives 121, 197–204. doi:10.1289/ehp.1104625

Hazus-MH, 2013. Hazus-MH 2.1 Technical Manual: Earthquake Model – Multi-hazard Loss Estimation Methodology. Department of Homeland Security, Federal Emergency Management Agency, Washington, DC, USA.

Heckerman, D., Horvitz, E., Middleton, B., 1993. An Approximate Nonmyopic Computation for Value of Information. IEEE Transactions on Pattern Recognition and Machine Intelligence 15, 292–298.

Heckerman, D.E., Horvitz, E.J., Nathwani, B.N., 1989. Update on the Pathfinder project, in: Proceedings of the Thirteenth Symposium on Computer Applications in Medical Care. pp. 203–207.

Hibbeler, R.C., 2009. Structural analysis, 7th ed. ed. Pearson/Prentice Hall, Upper Saddle River, N.J.

Hilton, R.W., 1977. The Determinants of Information Value: An Inquiry in the Context of the Cost-Volume-Profit Decision. Ohio State University.

Hondula, D.M., Barnett, A.G., 2014. Heat-Related Morbidity in Brisbane, Australia: Spatial Variation and Area-Level Predictors. Environmental Health Perspectives. doi:10.1289/ehp.1307496

Höppe, P., 1999. The physiological equivalent temperature - a universal index for the biometeorological assessment of the thermal environment. Int J Biometeorol 43, 71–75.

Hu, X., Daganzo, C., Madanat, S., 2015. A Reliability-Based Optimization Scheme for Maintenance Management in Large-Scale Bridge Networks, in: Proceedings of the 12th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP12). Vancouver, Canada.

Inaudi, D., 2011. Cost-Benefits Analysis in SHM Projects. Presented at the 5th International Conference on Structural Health Monitoring of Intelligent Infrastructure (SHMII-5), Cancún, México.

Inaudi, D., 2010. Optimal Design of Bridge SHM Systems Based on Risk and Opportunity Analysis, in: Proceedings of the Fifth International Conference on Bridge Maintenance, Safety and Management. Philadelphia, USA.

Inaudi, D., 2009. Overview of 40 Bridge Strucutral Health Monitoring Projects, in: Proceedings of the International Bridge Conference. Pittsburgh, USA.

Jensen, J.L.W.V., 1906. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. Acta Mathematica 30, 175–193. doi:10.1007/BF02418571

Kalkstein, L.S., Greene, J.S., 1997. An Evaluation of Climate/Mortality Relationships in Large U.S. Cities and the Possible Impacts of a Climate Change. Environmental Health Perspectives 105, 84–93.

Kisner, C., Mulder, K., VanGessel, B., 2012. Assessing Heat Vulnerability and Access to Cooling Centers in Detroit, Michigan. Presented at the IMAGIN 2013 Student Geospatial Poster & Paper Competition, University of Michigan, Ann Arbor, MI.

Kleiner, Y., 2001. Scheduling Inspection and Renewal of Large Infrastrucutre Assets. Journal of Infrastructure Systems 7, 136–143.

Koller, D., Friedman, N., 2009. Probabilistic graphical models: principles and techniques, Adaptive computation and machine learning. MIT Press, Cambridge, MA.

Konakli, K., Sudret, B., Faber, M.H., 2015. Numerical Investigations into the Value of Information in Lifecycle Analysis of Structural Systems. ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering B4015007. doi:10.1061/AJRUA6.0000850

Krause, A., 2008. Optimizing Sensing: Theory and Applications. Carnegie Mellon University, Pittsburgh, PA 15213.

Krause, A., Guestrin, C., 2009. Optimal Value of Information in Graphical Models. Journal of Artificial Intelligence Research 35, 557–591.

Krause, A., Guestrin, C., 2007. Nonmyopic Active Learning of Gaussian Processes: An Exploration–Exploitation Approach, in: Proceedings of the 24 Th International Conference on Machine Learning. Corvallis, Oregon, USA.

Krause, A., Leskovec, J., Guestrin, C., VanBriesen, J., Faloutsos, C., 2008a. Efficient sensor placement optimization for securing large water distribution networks. Journal of Water Resources Planning and Management 134, 516–526.

Krause, A., McMahon, H.B., Guestrin, C., Gupta, A., 2008b. Robust submodular observation selection. Journal of Machine Learning Research 9, 2761–2801.

Krause, A., Rajagopal, R., Gupta, A., Guestrin, C., 2011. Simultaneous Optimization of Sensor Placements and Balanced Schedules. IEEE Transactions on Automatic Control 56, 2390–2405. doi:10.1109/TAC.2011.2164010

Krause, A., Singh, A., Guestrin, C., 2008c. Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies. Journal of Machine Learning Research 9, 235–284.

Krause, A., Singh, A., Guestrin, C., 2008d. Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies. Journal of Machine Learning Research 9, 235–284.

Kröger, W., 2011. Vulnerable systems. Springer, London.

Kusaka, H., Kikegawa, Y., Kimura, F., Kondo, H., 2001. A Simple Single-Layer Urban Canopy Model For Atmospheric Models: Comparison With Multi-Layer And Slab Models. Boundary-Layer Meteorology 101, 329–358.

Lempert, R.J., Collins, M.T., 2007. Managing the Risk of Uncertain Threshold Responses: Comparison of Robust, Optimum, and Precautionary Approaches. Risk Analysis 27, 1009–1026. doi:10.1111/j.1539-6924.2007.00940.x

Li, D., Bou-Zeid, E., 2014. Quality and sensitivity of high-resolution numerical simulation of urban heat islands. Environmental Research Letters 9, 055001. doi:10.1088/1748-9326/9/5/055001

Li, D., Bou-Zeid, E., 2013. Synergistic Interactions between Urban Heat Islands and Heat Waves: The Impact in Cities Is Larger than the Sum of Its Parts*. Journal of Applied Meteorology and Climatology 52, 2051–2064. doi:10.1175/JAMC-D-13-02.1

Li, D., Bou-Zeid, E., Oppenheimer, M., 2014. The effectiveness of cool and green roofs as urban heat island mitigation strategies. Environmental Research Letters 9, 055002. doi:10.1088/1748-9326/9/5/055002

Lincke, T.R., 2002. Exploring the Computational Limits of Large Exhaustive Search Problems. Swiss Federal Institute of Technology, Zürich, Switzerland.

Lin-Jiong, Z., Yi-Min, L., Qing, B., Hai-Yang, Y., Guo-Xiong, W., 2012. Computational Performance of the High-Resolution Atmospheric Model FAMIL. Atmospheric and Oceanic Science Letters 5, 355–359. doi:10.1080/16742834.2012.11447024

Liu, M., Frangopol, D.M., 2006. Optimizing Bridge Network Maintenance Management under Uncertainty with Conflicting Criteria: Life-Cycle Maintenance, Failure, and User Costs. Journal of Structural Engineering 132, 1835–1845. doi:10.1061/(ASCE)0733-9445(2006)132:11(1835)

Loth, C., Baker, J.W., 2013. A spatial cross-correlation model of spectral accelerations at multiple periods. Earthquake Engineering & Structural Dynamics 42, 397–417. doi:10.1002/eqe.2212

Luque, J., Straub, D., 2016. Reliability analysis and updating of deteriorating systems with dynamic Bayesian networks. Structural Safety 62, 34–46. doi:10.1016/j.strusafe.2016.03.004

Malings, C., Memarzadeh, M., Pozzi, M., 2013. Optimal topology of sensor networks for management of infrastructure systems, in: Proceedings of the 6th International Conference on Structural Health Monitoring of Intelligent Infrastructure. Hong Kong, China.

Malings, C., Pozzi, M., 2016a. Conditional Entropy and Value of Information Metrics for Optimal Sensing in Infrastructure Systems. Structural Safety 60, 77–90.

Malings, C., Pozzi, M., 2016b. Value of Information for Spatially Distributed Systems: application to sensor placement. Reliability Engineering & System Safety 154, 219–233. doi:10.1016/j.ress.2016.05.010

Malings, C., Pozzi, M., 2016c. Value of information analysis for typical bridge network topologies, in: Proceedings of the 8th International Conference on Bridge Maintenance, Safety, and Management.

Malings, C., Pozzi, M., 2016d. Optimal sensing using value of information in spatio-temporal random fields. Presented at the 2016 working conference for the IFIP Working Group 7.5 on Reliability and Optimization of Structural Systems, Pittsburgh, PA, USA.

Malings, C., Pozzi, M., 2016e. Optimizing Sensing based on Decomposable Value of Information using Spatio-Temporal Probabilistic Models of Infrastructure Systems, in: Student Paper Competition of ASCE EMI SHMC Committee. Nashville, TN, USA.

Malings, C., Pozzi, M., 2015. Sensor Network Optimization using Bayesian Networks, Decision Graphs, and Value of Information, in: Proceedings of the 12th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP12). Vancouver, Canada.

Malings, C., Pozzi, M., 2014a. Conditional entropy & value of information metrics for informative sensing, in: Proceedings of the 17th Working Conference of the IFIP Working Group 7.5 on Reliability and Optimization of Structural Systems. Huangshan, PRC.

Malings, C., Pozzi, M., 2014b. Conditional entropy and value of information based metrics for sensor placement in infrastructure systems, in: Proceedings of the 6th World Conference on Structural Control and Monitoring. Barcelona, Spain.

Malings, C., Pozzi, M., Klima, K., Bou-Zeid, E., Ramamurthy, P., Bergés, M., 2017a. Surface Heat Assessment for Developed Environments: Probabilistic Urban Temperature Modeling; submitted to Computers, Environment, and Urban Systems.

Malings, C., Pozzi, M., Klima, K., Bou-Zeid, E., Ramamurthy, P., Bergés, M., 2017b. Surface Heat Assessment for Developed Environments: Optimizing Urban Temperature Monitoring; submitted to Computers, Environment, and Urban Systems.

Malings, C., Pozzi, M., Klima, K., Bou-Zeid, E., Ramamurthy, P., Bergés, M., 2016. Optimal Sensor Placement for Urban Heat Risk Response, in: 13th International Conference on Probabilistic Safety Assessment and Management (PSAM 13). Seoul, Republic of Korea.

Malings, C., Pozzi, M., Velibeyoglu, I., 2015. Sensor Placement Optimization for Structural Health Monitoring, in: Proceedings of the 10th International Workshop on Structural Health Monitoring.

Mascarenas, D., Plont, C., Brown, C., Cowell, M., Jameson, N.J., Block, J., Djidjev, S., Hahn, H., Farrar, C., 2014. A vibro-haptic human-machine interface for structural health monitoring. Structural Health Monitoring 13, 671–685. doi:10.1177/1475921714556569

Melillo, Richmond, T. C., Yohe, G. W., 2014. Climate Change Impacts in the United States: The Third National Climate Assessment. doi:10.7930/J0Z31WJ2

Meliou, A., Krause, A., Guestrin, C., Hellerstein, J.M., 2007. Nonmyopic Informative Path Planning in Spatio-Temporal Models, in: Proceedings of the 22nd National Conference on Artificial Intelligence (AAAI). Vancouver, Canada, pp. 602–607.

Mellor, G.L., Yamada, T., 1974. A Hierarchy of Turbulence Closure Models for Planetary Boundary Layers. Journal of the Atmospheric Sciences 31, 1791–1806. doi:10.1175/1520-0469(1974)031<1791:AHOTCM>2.0.CO;2

Memarzadeh, M., Pozzi, M., 2016. Value of information in sequential decision making: Component inspection, permanent monitoring and system-level scheduling. Reliability Engineering & System Safety 154, 137–151. doi:10.1016/j.ress.2016.05.014

Memarzadeh, M., Pozzi, M., 2015. Integrated Inspection Scheduling and Maintenance Planning for Infrastructure Systems: Integrated inspection scheduling and maintenance planning. Computer-Aided Civil and Infrastructure Engineering n/a-n/a. doi:10.1111/mice.12178

Memarzadeh, M., Pozzi, M., Kolter, J.Z., 2016. Hierarchical modeling of systems with similar components: A framework for adaptive monitoring and control. Reliability Engineering & System Safety 154, 137–151. doi:doi:10.1016/j.ress.2016.05.014

Memarzadeh, M., Pozzi, M., Zico Kolter, J., 2015. Optimal Planning and Learning in Uncertain Environments for the Management of Wind Farms. Journal of Computing in Civil Engineering 29, 04014076. doi:10.1061/(ASCE)CP.1943-5487.0000390

Menne, B., Matthies, F., 2009. EuroHEAT: Improving public health responses to extreme weather/heat-waves, technical summary. World Health Organization.

Micron Optics, 2012. Optical Sensing Instrumentation & Software - Enlight User Guide (No. 1.138). Micron Optics, Inc.

Min, H.-S.J., Beyeler, W.E., Brown, T.J., Son, Y.J., Jones, A.T., 2005. Toward Modeling and Simulation of Critical National Infrastructure Interdependencies (No. SAND2005–4828J). Sandia National Laboratory.

Mlawer, E.J., Taubman, S.J., Brown, P.D., Iacono, M.J., Clough, S.A., 1997. Radiative transfer for inhomogeneous atmospheres: RRTM, a validated correlated-k model for the longwave. Journal of Geophysical Research: Atmospheres 102, 16663–16682. doi:10.1029/97JD00237

Mokrech, M., Nicholls, R.J., Dawson, R.J., 2012. Scenarios of Future Built Environment for Coastal Risk Assessment of Climate Change Using a GIS-Based Multicriteria Analysis. Environment and Planning B: Planning and Design 39, 120–136. doi:10.1068/b36077

Morimoto, R., 2010. Estimating the benefits of effectively and proactively maintaining infrastructure with the innovative Smart Infrastructure sensor system. Socio-Economic Planning Sciences 44, 247–257. doi:10.1016/j.seps.2010.07.005

Nemhauser, G., Wolsey, L., Fisher, M., 1978. An analysis of the approximations for maximizing submodular set functions. Mathematical Programming 14, 265–294.

New York City Panel on Climate Change, 2013. Climate Risk Information 2013: Observations, Climate Change Projections, and Maps, NPCC2. Prepared for use by the City of New York Special Initiative on Rebuilding and Resiliency, New York, New York.

Newman, M.E.J., 2010. Networks: an introduction. Oxford University Press, Oxford ; New York.

Oakley, J.E., Brennan, A., Tappenden, P., Chilcott, J., 2010. Simulation sample sizes for Monte Carlo partial EVPI calculations. Journal of Health Economics 29, 468–477. doi:10.1016/j.jhealeco.2010.03.006

Oke, T.R., 1997. Urban Climates and Global Environmental Change, in: Applied Climatology: Principles & Practices. Routledge, New York, NY, pp. 273–287.

Papadimitriou, C., 2004. Optimal sensor placement methodology for parametric identification of structural systems. Journal of Sound and Vibration 278, 923–947. doi:10.1016/j.jsv.2003.10.063

Papaioannou, I., Betz, W., Zwirglmaier, K., Straub, D., 2015. MCMC algorithms for Subset Simulation. Probabilistic Engineering Mechanics 41, 89–103. doi:10.1016/j.probengmech.2015.06.006

Parandehgheibi, M., Modiano, E., 2013. Robustness of interdependent networks: The case of communication networks and the power grid. IEEE, pp. 2164–2169. doi:10.1109/GLOCOM.2013.6831395

Parson, M., Jameson, J., 2012. Pittsburgh Climate Action Plan, version 2.0. Pittsburgh Climate Initiative, Green Building Alliance, Pittsburgh, PA.

Pozzi, M., Der Kiureghian, A., 2011a. Computation of lifetime value of information for monitoring systems, in: Proceedings of the 8th International Workshop on Structural Health Monitoring. Presented at the 8th International Workshop on Structural Health Monitoring, Stanford University, California, USA.

Pozzi, M., Der Kiureghian, A., 2011b. Assessing the value of information for long-term structural health monitoring, in: Kundu, T. (Ed.), Proceedings of SPIE, the International Society of Photo-Optical Instrumentation Engineers. p. 79842W–79842W–14. doi:10.1117/12.881918

Qin, J., Thöns, S., Faber, M.H., 2015. On the value of SHM in the context of service life integrity management, in: Proceedings of the 12th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP12). Vancouver, Canada.

Raiffa, H., Schlaifer, R., 1961. Applied Statistical Decision Theory. Harvard University Press, Cambridge, Massachusetts, USA.

Ramamurthy, P., Li, D., Bou-Zeid, E., 2015. High-resolution simulation of heatwave events in New York City. Theoretical and Applied Climatology. doi:10.1007/s00704-015-1703-8

Rasmussen, C.E., Williams, C.K.I., 2006. Gaussian processes for machine learning, Adaptive computation and machine learning. MIT Press, Cambridge, Mass.

Reid, C., O'Neill, M., Gronlund, C., Brines, S., Brown, D., Diez-Roux, A., Schwartz, J., 2009. Mapping Community Determinants of Heat Vulnerability. Environmental Health Perspectives. doi:10.1289/ehp.0900683

Rhoades, E.K., Contreras, C., Garrett, S.K., Bakshi, M., Bellomo, A.J., 2014. Your Health and Climate Change in Los Angeles County (No. 1), Climate and Health Series. Los Angeles County Department of Public Health.

Rinaldi, S.M., Peerenboom, J.P., Kelly, T.K., 2001. Identifying, understanding, and analyzing critical infrastructure interdependencies. IEEE Control Systems Magazine 21, 11–25. doi:10.1109/37.969131

Robine, J.-M., Cheung, S.L.K., Le Roy, S., Van Oyen, H., Griffiths, C., Michel, J.-P., Herrmann, F.R., 2008. Death toll exceeded 70,000 in Europe during the summer of 2003. Comptes Rendus Biologies 331, 171–178. doi:10.1016/j.crvi.2007.12.001

Rosenhead, J., 2013. Robustness Analysis, in: Gass, S.I., Fu, M.C. (Eds.), Encyclopedia of Operations Research and Management Science. Springer US, Boston, MA, pp. 1346–1347. doi:10.1007/978-1-4419-1153-7_200722

Rosenthal, J.K., Kinney, P.L., Metzger, K.B., 2014. Intra-urban vulnerability to heat-related mortality in New York City, 1997–2006. Health & Place 30, 45–60. doi:10.1016/j.healthplace.2014.07.014

Sailor, D.J., 2002. Urban Heat Islands, Opportunities and Challenges for Mitigation and Adaptation.

Sankar, S., 2012. The rise of human-computer cooperation, TED.

Satumtira, G., Dueñas-Osorio, L., 2010. Synthesis of Modeling and Simulation Methods on Critical Infrastructure Interdependencies Research, in: Gopalakrishnan, K., Peeta, S. (Eds.), Sustainable and Resilient Critical Infrastructure Systems. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 1–51.

Schrijver, A., 2013. A Course in Combinatorial Optimization. Department of Mathematics, University of Amsterdam, The Netherlands.

Schweckendiek, T., Vrouwenvelder, T., 2015. Value of information in retrofitting of flood defenses, in: Proceedings of the 12th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP12). Vancouver, Canada.

Sheth, A., 2013. Transforming Big Data into Smart Data: Deriving Value via Harnessing Volume, Variety, and Velocity Using Semantics and Semantic Web, in: Keynote Talk Presented at the 21st Italian Symposium on Advanced Database Systems, Roccella Jonica, Italy.

Simon, H.A., 1959. Theories of decision-making in economics and behavioral science. The American economic review 49, 253–283.

Singhvi, V., Krause, A., Guestrin, C., Garrett, Jr., J.H., Matthews, H.S., 2005. Intelligent Light Control using Sensor Networks, in: Proceedings of the 3rd ACM Conference on Embedded Networked Sensor Systems (SenSys'05). Presented at the 3rd ACM Conference on Embedded Networked Sensor Systems (SenSys'05), San Diego, California, USA.

Skamarock, W.C., Klemp, J.B., Dudhia, J., Gill, D.O., Barker, D.M., Wang, W., Powers, J.G., 2005. A Description of the Advanced Research WRF Version 2 (No. NCAR/TN–468+STR). National Center for Atmospheric Research.

Špačková, O., Straub, D., 2015. Cost-Benefit Analysis for Optimization of Risk Protection Under Budget Constraints: CBA for Optimization of Risk Protection. Risk Analysis 35, 941–959. doi:10.1111/risa.12310

Sper de Almeida, E., Bauer, M., 2012. Reducing time delays in computing numerical weather models at regional and local levels: a grid-based approach. International Journal of Grid Computing & Applications 3, 1–17.

Stafoggia, M., Forastiere, F., Agostini, D., Caranci, N., de'Donato, F., Demaria, M., Michelozzi, P., Miglio, R., Rognoni, M., Russo, A., Perucci, C.A., 2008. Factors affecting in-hospital heat-related mortality: a multi-city case-crossover analysis. Journal of Epidemiology & Community Health 62, 209–215. doi:10.1136/jech.2007.060715

Stenström, C., Norrbin, P., Parida, A., Kumar, U., 2016. Preventive and corrective maintenance – cost comparison and cost–benefit analysis. Structure and Infrastructure Engineering 12, 603–617. doi:10.1080/15732479.2015.1032983

Straub, D., 2014. Value of Information Analysis with Structural Reliability Methods. Structural Safety 49, 75–86.

Straub, D., 2004. Generic Approaches to Risk Based Inspection Planning for Steel Structures. ETH Zürich, Zürich, Switzerland.

Straub, D., Der Kiureghian, A., 2010a. Bayesian Network Enhanced with Structural Reliability Methods: Application. Journal of Engineering Mechanics 136, 1259–1270. doi:10.1061/(ASCE)EM.1943-7889.0000170

Straub, D., Der Kiureghian, A., 2010b. Bayesian Network Enhanced with Structural Reliability Methods: Methodology. Journal of Engineering Mechanics 136, 1248–1258. doi:10.1061/(ASCE)EM.1943-7889.0000173

Straub, D., Faber, M.H., 2006. Computational Aspects of Risk-Based Inspection Planning. Computer-Aided Civil and Infrastructure Engineering 21, 179–192.

Straub, D., Faber, M.H., 2005. Risk based inspection planning for structural systems. Structural Safety 27, 335–355. doi:10.1016/j.strusafe.2005.04.001

Straub, D., Faber, M.H., 2004. System effects in generic risk based inspection planning. ASME Journal of Offshore Mechanics and Arctic Engineering 126, 265–271.

Sui, Y., Gotovos, A., Burdick, J.W., Krause, A., 2015. Safe Exploration for Optimization with Gaussian Processes, in: Proceedings of the 32nd International Conference on Machine Learning. Presented at the 32nd International Conference on Machine Learning, Lille, France.

Sustainable DC, 2013. Sustainable DC Plan. Washington DC.

Swanson, L., 2001. Linking maintenance strategies to performance. International Journal of Production Economics 70, 237–244. doi:10.1016/S0925-5273(00)00067-0

Takahashi, Y., Der Kiureghian, A., Ang, A.H.-S., 2004. Life-cycle cost analysis based on a renewal model of earthquake occurrences. Earthquake Engineering and Structural Dynamics 33, 859–880.

Thöns, S., Schneider, R., Faber, M.H., 2015. Quantification of the value of structural health monitoring information for fatigue deteriorating structural systems, in: Proceedings of the 12th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP12). Vancouver, Canada.

US EPA, 2012. Reducing Urban Heat Island: Compendium of Strategies - Urban Heat Island Basics. U.S. Environmental Protection Agency, Washington, DC, USA.

Vanchinathan, H.P., Marfurt, A., Robelin, C.-A., Kossmann, D., Krause, A., 2015. Discovering Valuable items from Massive Data. ACM Press, pp. 1195–1204. doi:10.1145/2783258.2783360

Wang, Z.-H., Bou-Zeid, E., Smith, J.A., 2013. A coupled energy transport and hydrological model for urban canopies evaluated using a wireless sensor network. Quarterly Journal of the Royal Meteorological Society 139, 1643–1657. doi:10.1002/qj.2032

Wells, E., Klima, K., 2016. Excessive Heat Event Risk & Response: How the Public and Public Health Experts Prepare, Respond, and Interpret Extreme Heat. Cities and the Environment In Review.

Worley, P.H., Mirin, A.A., Craig, A.P., Taylor, M.A., Dennis, J.M., Vertenstein, M., 2011. Performance of the community earth system model. ACM Press, p. 1. doi:10.1145/2063384.2063457

Yi, T.-H., Li, H.-N., 2012. Methodology Developments in Sensor Placement for Health Monitoring of Civil Infrastructures. International Journal of Distributed Sensor Networks 2012, 11. doi:10.1155/2012/612726

Yi, T.-H., Li, H.-N., Gu, M., 2011. Optimal sensor placement for structural health monitoring based on multiple optimization strategies. The Structural Design of Tall and Special Buildings 20, 881–900. doi:10.1002/tal.712

Yoshida, I., 2015. Parameter study on optimal sampling planning based on value of information, in: Proceedings of the 12th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP12). Vancouver, Canada.

Yuen, K.-V., Kuok, S.-C., 2015. Efficient Bayesian sensor placement algorithm for structural identification: a general approach for multi-type sensory systems. Earthquake Engineering & Structural Dynamics 44, 757–774. doi:10.1002/eqe.2486

Zelaschi, C., Pozzi, M., Monteiro, R., Malings, C., 2017. System-Level Fragility Assessment under Interdependent Spatial Demand and Structural Capacity. Presented at the 16th World Conference on Earthquake Engineering, Santiago, Chile.

Appendix A: Value of Information Bounds and Relationships

This appendix presents derivations supporting the assertions in Section 4.1 concerning the upper and lower bounds on VoI, as well as the equivalence of the regret-based and loss-based formulations for the VoI and the relative performance of online and offline optimization approaches for the VoI metric. The basic modeling assumptions of Section 3.3 are relaxed in this appendix.

A.1. Lower Bound for the Value of Information

This section derives a lower-bound for the VoI. This derivation is adapted from the work of Hilton (1977). Begin by noting that the following expression holds true for all $\mathbf{y} \in \mathrm{Dom}(Y)$:

$$\min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}}\mathrm{L}(\mathbf{f},\mathbf{a}) \leq \mathbb{E}_{F|\mathbf{y}}\mathrm{L}(\mathbf{f},\mathbf{a})$$

This is a consequence of the definition of the minimization, i.e., the minimum over $\mathbf{a}$ of $\mathbb{E}_{F|\mathbf{y}}\mathrm{L}(\mathbf{f},\mathbf{a})$, a function of $\mathbf{y}$ alone, is by definition a lower bound for $\mathbb{E}_{F|\mathbf{y}}\mathrm{L}(\mathbf{f},\mathbf{a})$, which is a function of $\mathbf{a}$ and $\mathbf{y}$. Since this relationship holds for all $\mathbf{y} \in \mathrm{Dom}(Y)$, it will hold under the expectation over all $\mathbf{y}$ as well. Taking the expectation of both sides with respect to Y yields:

$$\mathbb{E}_Y \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}}\mathrm{L}(\mathbf{f},\mathbf{a}) \leq \mathbb{E}_Y\mathbb{E}_{F|\mathbf{y}}\mathrm{L}(\mathbf{f},\mathbf{a})$$

Note that the left-hand side is the definition of the posterior expected loss from Equation 4-5. On the right-hand side, the expectations can be replaced with a joint expectation over F and Y:

$$\mathbb{E}\mathrm{L}(Y) \leq \mathbb{E}_{F,Y}\mathrm{L}(\mathbf{f},\mathbf{a})$$

Since $\mathrm{L}(\mathbf{f},\mathbf{a})$ is not a function of $\mathbf{y}$, this joint expectation is equivalent to an expectation over F only. Taking the minimum of both sides with respect to A preserves with inequality, since it must hold over all $\mathbf{a} \in \mathcal{A}$:

$$\min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}\mathrm{L}(Y) \leq \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_F\mathrm{L}(\mathbf{f},\mathbf{a})$$

On the left-hand side, $\mathbb{E}\mathrm{L}(Y)$ is not a function of $\mathbf{a}$, so it can be replaced with simply $\mathbb{E}\mathrm{L}(Y)$. The right-hand side is the definition for the prior expected loss from Equation 4-1:

$$\mathbb{E}\mathrm{L}(Y) \leq \mathbb{E}\mathrm{L}(\emptyset)$$

Subtracting $\mathbb{E}\mathrm{L}(Y)$ from both sides proves the bound:

$$0 \leq \mathbb{E}\mathrm{L}(\emptyset) - \mathbb{E}\mathrm{L}(Y) = \mathrm{VoI}(Y)$$

A.2. Upper Bound for the Value of Information

The following proof makes use of Jensen's Inequality for concave functions. Concavity and Jensen's Inequality are common concepts in mathematics and statistics, discussed by, e.g., Casella and Berger (2002, pp. 189–190).

The first step towards proving an upper-bound for VoI is to prove the concavity of the minimization function. Concavity can be proved by showing that the following expression holds true for any $\alpha \in [0,1]$:

$$\min_{\mathbf{a}\in\mathcal{A}}[\alpha\mathrm{L}(\mathbf{f}_1,\mathbf{a}) + (1-\alpha)\mathrm{L}(\mathbf{f}_2,\mathbf{a})] \geq \alpha \min_{\mathbf{a}\in\mathcal{A}} \mathrm{L}(\mathbf{f}_1,\mathbf{a}) + (1-\alpha) \min_{\mathbf{a}\in\mathcal{A}} \mathrm{L}(\mathbf{f}_2,\mathbf{a})$$

Let $\mathbf{a}_1^*$ be the action set that minimizes $\mathrm{L}(\mathbf{f}_1,\mathbf{a})$, $\mathbf{a}_2^*$ be the action set that minimizes $\mathrm{L}(\mathbf{f}_2,\mathbf{a})$, and $\mathbf{a}_{\mathrm{sum}}^*$ be the action set that minimizes the sum $\alpha\mathrm{L}(\mathbf{f}_1,\mathbf{a}) + (1-\alpha)\mathrm{L}(\mathbf{f}_2,\mathbf{a})$. By definition of the minimum:

$$\mathrm{L}(\mathbf{f}_1,\mathbf{a}_{\mathrm{sum}}^*) \geq \mathrm{L}(\mathbf{f}_1,\mathbf{a}_1^*)$$

and:

$$\mathrm{L}(\mathbf{f}_2,\mathbf{a}_{\mathrm{sum}}^*) \geq \mathrm{L}(\mathbf{f}_2,\mathbf{a}_2^*)$$

Multiplying these two inequalities by positive constants (α and $1-\alpha$, respectively) preserves the inequality, as does adding them together:

$$\alpha\mathrm{L}(\mathbf{f}_1,\mathbf{a}_{\mathrm{sum}}^*) + (1-\alpha)\mathrm{L}(\mathbf{f}_2,\mathbf{a}_{\mathrm{sum}}^*) \geq \alpha\mathrm{L}(\mathbf{f}_1,\mathbf{a}_1^*) + (1-\alpha)\mathrm{L}(\mathbf{f}_2,\mathbf{a}_2^*)$$

The left-hand side of the inequality can be expressed as:

$$\alpha\mathrm{L}(\mathbf{f}_1,\mathbf{a}_{\mathrm{sum}}^*) + (1-\alpha)\mathrm{L}(\mathbf{f}_2,\mathbf{a}_{\mathrm{sum}}^*) = \min_{\mathbf{a}\in\mathcal{A}}[\alpha\mathrm{L}(\mathbf{f}_1,\mathbf{a}) + (1-\alpha)\mathrm{L}(\mathbf{f}_2,\mathbf{a})]$$

while the right-hand side can be expressed as:

$$\alpha\mathrm{L}(\mathbf{f}_1,\mathbf{a}_1^*) + (1-\alpha)\mathrm{L}(\mathbf{f}_2,\mathbf{a}_2^*) = \alpha \min_{\mathbf{a}\in\mathcal{A}} \mathrm{L}(\mathbf{f}_1,\mathbf{a}) + (1-\alpha) \min_{\mathbf{a}\in\mathcal{A}} \mathrm{L}(\mathbf{f}_2,\mathbf{a})$$

This is the original inequality, proving the concavity of the minimization function.

Concavity for the minimization function means Jensen's Inequality for concave functions will hold true for the minimization, i.e., that the minimum of the expected value of the loss function over F is greater than or equal to the expectation over F of the minimum loss:

$$\min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_F\mathrm{L}(\mathbf{f},\mathbf{a}) \geq \mathbb{E}_F \min_{\mathbf{a}\in\mathcal{A}} \mathrm{L}(\mathbf{f},\mathbf{a})$$

This inequality will hold even if the expectation is conditioned on another variable Y (of which the loss is not a function):

$$\min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}}\mathrm{L}(\mathbf{f},\mathbf{a}) \geq \mathbb{E}_{F|\mathbf{y}} \min_{\mathbf{a}\in\mathcal{A}} \mathrm{L}(\mathbf{f},\mathbf{a})$$

As this inequality holds in general for any $\mathbf{y} \in \mathrm{Dom}(Y)$, it will hold after the expectation over Y as well:

$$\mathbb{E}_Y \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}}\mathrm{L}(\mathbf{f},\mathbf{a}) \geq \mathbb{E}_Y\mathbb{E}_{F|\mathbf{y}} \min_{\mathbf{a}\in\mathcal{A}} \mathrm{L}(\mathbf{f},\mathbf{a})$$

For the expression on the right-hand side, the two expectations can be combined into an expectation over F and Y using the chain rule:

$$\mathbb{E}_Y \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}}\mathrm{L}(\mathbf{f},\mathbf{a}) \geq \mathbb{E}_{F,Y} \min_{\mathbf{a}\in\mathcal{A}} \mathrm{L}(\mathbf{f},\mathbf{a})$$

Furthermore, since $\min_{\mathbf{a}\in\mathcal{A}} \mathrm{L}(\mathbf{f},\mathbf{a})$ is not a function of $\mathbf{y}$, the joint expectation can be replaced with an expectation over F only:

$$\mathbb{E}_Y \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}}\mathrm{L}(\mathbf{f},\mathbf{a}) \geq \mathbb{E}_F \min_{\mathbf{a}\in\mathcal{A}} \mathrm{L}(\mathbf{f},\mathbf{a})$$

Finally, noting that $\mathbb{E}_{F|\mathbf{f}}\mathrm{L}(\mathbf{f},\mathbf{a})$ is equivalent to $\mathrm{L}(\mathbf{f},\mathbf{a})$ (taking the expectation of a function over a fixed quantity is equivalent to substituting in that quantity to the function) gives:

$$\mathbb{E}_Y \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}}\mathrm{L}(\mathbf{f},\mathbf{a}) \geq \mathbb{E}_F \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{f}}\mathrm{L}(\mathbf{f},\mathbf{a})$$

Both sides of the inequality now have the form of Equation 4-5, where the left-hand side is the posterior expected loss given an observation of Y and the right-hand side is the posterior expected loss given an observation of F:

$$\mathbb{E}\mathrm{L}(Y) \geq \mathbb{E}\mathrm{L}(F)$$

Multiplying both sides by -1 reverses the inequality, while adding $\mathbb{E}\mathrm{L}(\emptyset)$ to both sides preserves it:

$$\mathbb{E}\mathrm{L}(\emptyset) - \mathbb{E}\mathrm{L}(Y) \leq \mathbb{E}\mathrm{L}(\emptyset) - \mathbb{E}\mathrm{L}(F)$$

This gives the upper-bound for $\mathrm{VoI}(Y)$:

$$\mathrm{VoI}(Y) \leq \mathrm{VoI}(F)$$

A.3. Regret-based and Loss-based Formulation Equivalence

This section provides a proof for the equivalence of the loss-based and regret-based formulations of VoI in Equations 4-6 and 4-11. Substituting the definition of regret from Equation 4-8 into the prior expected regret formula of Equation 4-10 yields:

$$\mathbb{E}\mathrm{R}(\emptyset) = \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_F\left[\mathrm{L}(\mathbf{f},\mathbf{a}) - \mathrm{L}\left(\mathbf{f},\mathbf{a}^*(\mathbf{f})\right)\right]$$

Using the linearity of the expectation, as well as the fact that the second half of the expression in brackets is not a function of $\mathbf{a}$, the prior expected regret can be expressed as:

$$\mathbb{E}\mathrm{R}(\emptyset) = \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_F\mathrm{L}(\mathbf{f},\mathbf{a}) - \mathbb{E}_F\left[\min_{\mathbf{a}\in\mathcal{A}} \mathrm{L}(\mathbf{f},\mathbf{a})\right]$$

where the definition of $\mathbf{a}^*(\mathbf{f})$ from Equation 4-9 is substituted into the second half of the expression. Because the expectation over F conditioned on $\mathbf{f}$ of $\mathrm{L}(\mathbf{f},\mathbf{a})$ is equivalent to $\mathrm{L}(\mathbf{f},\mathbf{a})$ directly (taking the expectation of a function over a fixed quantity is equivalent to substituting in that quantity to the function), the prior expected regret is expressed equivalently as:

$$\mathbb{E}\mathrm{R}(\emptyset) = \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_F\mathrm{L}(\mathbf{f},\mathbf{a}) - \mathbb{E}_F\left[\min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{f}}\,\mathrm{L}(\mathbf{f},\mathbf{a})\right] = \mathbb{E}\mathrm{L}(\emptyset) - \mathbb{E}\mathrm{L}(F) = \mathrm{VoI}(F)$$

Note therefore that the prior expected regret is equivalent to the value of complete perfect information $\mathrm{VoI}(F)$ on the random field.

Similarly, for the posterior expected regret of Equation 4-10:

$$\mathbb{E}\mathrm{R}(Y) = \mathbb{E}_Y \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}}\mathrm{L}(\mathbf{f},\mathbf{a}) - \mathbb{E}_Y\mathbb{E}_{F|\mathbf{y}}\left[\min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{f}}\,\mathrm{L}(\mathbf{f},\mathbf{a})\right] = \mathbb{E}\mathrm{L}(Y) - \mathbb{E}\mathrm{L}(F)$$

Therefore, the value of information definitions in Equations 4-6 and 4-11 are equivalent:

$$\mathbb{E}\mathrm{L}(\emptyset) - \mathbb{E}\mathrm{L}(Y) = \left(\mathbb{E}\mathrm{L}(\emptyset) - \mathbb{E}\mathrm{L}(F)\right) - \left(\mathbb{E}\mathrm{L}(Y) - \mathbb{E}\mathrm{L}(F)\right) = \mathbb{E}\mathrm{R}(\emptyset) - \mathbb{E}\mathrm{R}(Y)$$

A.4. Greedy Online versus Offline Sensor Placement

This section demonstrates that greedy online sensor placement using the VoI metric will outperform greedy offline sensor placement. As a preliminary step, the following relationship is proven:

$$\mathbb{E}_{\left\{Y_1^*,\dots,Y_{j-1}^*\right\}}\mathrm{M}_{\mathrm{VoI}}\left(Y_j|\mathbf{y}_{j-1}\right) = \mathrm{M}_{\mathrm{VoI}}\left(\left\{Y_1^*,\dots,Y_{j-1}^*,Y_j\right\}\right) - \mathrm{M}_{\mathrm{VoI}}\left(\left\{Y_1^*,\dots,Y_{j-1}^*\right\}\right)$$

That is, the expected value of the VoI metric for online optimization in Equation 4-13 over potential observations for previously placed measurements is equal to the difference of the values of the offline sensor placement objective of Equation 4-12 with and without the next

measurement Y_j. To show that this relationship holds, first substitute in the definition of Equation 4-13 in the left-hand side and Equation 4-12 in the right-hand side:

$$\begin{aligned}\mathbb{E}_{\{Y_1^*,\dots,Y_{j-1}^*\}}\left[\mathrm{VoI}\left(Y_j|\mathbf{y}_{j-1}\right) - \mathrm{C}\left(\{Y_1^*,\dots,Y_{j-1}^*,Y_j\}\right) + \mathrm{C}\left(\{Y_1^*,\dots,Y_{j-1}^*\}\right)\right] \\ = \mathrm{VoI}\left(\{Y_1^*,\dots,Y_{j-1}^*,Y_j\}\right) - \mathrm{C}\left(\{Y_1^*,\dots,Y_{j-1}^*,Y_j\}\right) - \mathrm{VoI}\left(\{Y_1^*,\dots,Y_{j-1}^*\}\right) \\ + \mathrm{C}\left(\{Y_1^*,\dots,Y_{j-1}^*\}\right)\end{aligned}$$

Note that the cost functions are constant with respect to the values of the measurements, and so they cancel each other on both sides:

$$\mathbb{E}_{\{Y_1^*,\dots,Y_{j-1}^*\}}\left[\mathrm{VoI}\left(Y_j|\mathbf{y}_{j-1}\right)\right] = \mathrm{VoI}\left(\{Y_1^*,\dots,Y_{j-1}^*,Y_j\}\right) - \mathrm{VoI}\left(\{Y_1^*,\dots,Y_{j-1}^*\}\right)$$

Substituting the definition of VoI from Equation 4-6 in the right-hand side:

$$\mathbb{E}_{\{Y_1^*,\dots,Y_{j-1}^*\}}\left[\mathrm{VoI}\left(Y_j|\mathbf{y}_{j-1}\right)\right] = \mathbb{E}\mathrm{L}(\emptyset) - \mathbb{E}\mathrm{L}\left(\{Y_1^*,\dots,Y_{j-1}^*,Y_j\}\right) - \mathbb{E}\mathrm{L}(\emptyset) + \mathbb{E}\mathrm{L}\left(\{Y_1^*,\dots,Y_{j-1}^*\}\right)$$

where the prior expected losses cancel each other out. Substituting the definition from Equation 4-14 on the left-hand side:

$$\begin{aligned}\mathbb{E}_{\{Y_1^*,\dots,Y_{j-1}^*\}}\left[\min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}_{j-1}}\mathrm{L}(\mathbf{f},\mathbf{a}) - \mathbb{E}_{Y_j|\mathbf{y}_{j-1}} \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\{\mathbf{y}_{j-1},y_j\}}\mathrm{L}(\mathbf{f},\mathbf{a})\right] \\ = \mathbb{E}\mathrm{L}\left(\{Y_1^*,\dots,Y_{j-1}^*\}\right) - \mathbb{E}\mathrm{L}\left(\{Y_1^*,\dots,Y_{j-1}^*,Y_j\}\right)\end{aligned}$$

Using the linearity of the expectation, this can be applied to the two terms on the left-hand side separately. For the second term on the left-hand side, the chain rule is used to replace $\mathbb{E}_{\{Y_1^*,\dots,Y_{j-1}^*\}}\mathbb{E}_{Y_j|\mathbf{y}_{j-1}}$ with $\mathbb{E}_{\{Y_1^*,\dots,Y_{j-1}^*,Y_j\}}$.

$$\begin{aligned}\mathbb{E}_{\{Y_1^*,\dots,Y_{j-1}^*\}} \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}_{j-1}}\mathrm{L}(\mathbf{f},\mathbf{a}) - \mathbb{E}_{\{Y_1^*,\dots,Y_{j-1}^*,Y_j\}} \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\{\mathbf{y}_{j-1},y_j\}}\mathrm{L}(\mathbf{f},\mathbf{a}) \\ = \mathbb{E}\mathrm{L}\left(\{Y_1^*,\dots,Y_{j-1}^*\}\right) - \mathbb{E}\mathrm{L}\left(\{Y_1^*,\dots,Y_{j-1}^*,Y_j\}\right)\end{aligned}$$

The two terms on the left-hand side are now the definitions of the posterior expected loss from Equation 4-5 for sets $\{Y_1^*,\dots,Y_{j-1}^*\}$ and $\{Y_1^*,\dots,Y_{j-1}^*,Y_j\}$ respectively. Substituting these definitions shows that both sides are equal, proving the relationship:

$$\mathbb{E}\mathrm{L}\left(\{Y_1^*,\dots,Y_{j-1}^*\}\right) - \mathbb{E}\mathrm{L}\left(\{Y_1^*,\dots,Y_{j-1}^*,Y_j\}\right) = \mathbb{E}\mathrm{L}\left(\{Y_1^*,\dots,Y_{j-1}^*\}\right) - \mathbb{E}\mathrm{L}\left(\{Y_1^*,\dots,Y_{j-1}^*,Y_j\}\right) \blacksquare$$

Using this relationship, the relative performance of greedy online and offline sensor placement is investigated. First, by definition of the maximum, the following relationship holds true for all $Y_j \in \mathcal{Y}_{j-1}$ and for all potential values of $\mathbf{y}_{j-1}$:

$$\max_{Y_j \in \mathcal{Y}_{j-1}} \mathrm{M}_{\mathrm{VoI}}\left(Y_j | \mathbf{y}_{j-1}\right) \geq \mathrm{M}_{\mathrm{VoI}}\left(Y_j | \mathbf{y}_{j-1}\right)$$

Since this holds true for any value of $\mathbf{y}_{j-1}$, taking the expectation over possible values preserves the inequality:

$$\mathbb{E}_{\left\{Y_1^*, \ldots, Y_{j-1}^*\right\}} \max_{Y_j \in \mathcal{Y}_{j-1}} \mathrm{M}_{\mathrm{VoI}}\left(Y_j | \mathbf{y}_{j-1}\right) \geq \mathbb{E}_{\left\{Y_1^*, \ldots, Y_{j-1}^*\right\}} \mathrm{M}_{\mathrm{VoI}}\left(Y_j | \mathbf{y}_{j-1}\right)$$

Substituting the relationship shown at the beginning of this section into the right-hand side yields:

$$\mathbb{E}_{\left\{Y_1^*, \ldots, Y_{j-1}^*\right\}} \max_{Y_j \in \mathcal{Y}_{j-1}} \mathrm{M}_{\mathrm{VoI}}\left(Y_j | \mathbf{y}_{j-1}\right) \geq \mathrm{M}_{\mathrm{VoI}}\left(\left\{Y_1^*, \ldots, Y_{j-1}^*, Y_j\right\}\right) - \mathrm{M}_{\mathrm{VoI}}\left(\left\{Y_1^*, \ldots, Y_{j-1}^*\right\}\right)$$

As this relationship holds for any choice of $Y_j \in \mathcal{Y}_{j-1}$, it will hold for the Y_j that maximizes the right-hand side:

$$\mathbb{E}_{\left\{Y_1^*, \ldots, Y_{j-1}^*\right\}} \max_{Y_j \in \mathcal{Y}_{j-1}} \mathrm{M}_{\mathrm{VoI}}\left(Y_j | \mathbf{y}_{j-1}\right) \geq \max_{Y_j \in \mathcal{Y}_{j-1}} \left[\mathrm{M}_{\mathrm{VoI}}\left(\left\{Y_1^*, \ldots, Y_{j-1}^*, Y_j\right\}\right) - \mathrm{M}_{\mathrm{VoI}}\left(\left\{Y_1^*, \ldots, Y_{j-1}^*\right\}\right)\right]$$

Since the second term on the right-hand side is not a function of Y_j, it can be removed from the maximization:

$$\mathbb{E}_{\left\{Y_1^*, \ldots, Y_{j-1}^*\right\}} \max_{Y_j \in \mathcal{Y}_{j-1}} \mathrm{M}_{\mathrm{VoI}}\left(Y_j | \mathbf{y}_{j-1}\right) \geq \max_{Y_j \in \mathcal{Y}_{j-1}} \mathrm{M}_{\mathrm{VoI}}\left(\left\{Y_1^*, \ldots, Y_{j-1}^*, Y_j\right\}\right) - \mathrm{M}_{\mathrm{VoI}}\left(\left\{Y_1^*, \ldots, Y_{j-1}^*\right\}\right)$$

Note that the left-hand side is the expected value, over potential observations of previous measurements, of the optimal value of the online greedy optimization objective of Equation 2-18, which is the incremental benefit in terms of net VoI of adding the next greedily selected sensor to the set. The right-hand side is the difference of the optimal value of the offline greedy optimization objective of Equation 2-17 and the offline objective value for previously selected measurements. In other words, the left-hand side is the expected gain in terms of the VoI metric for the next greedily selected online measurement, and the right-hand side is the expected gain for the next greedily selected offline measurement. This relationship shows that, in an expected sense, the net VoI gain from greedy online sensor placement exceeds that from greedy offline sensor placement.

Appendix B: Submodularity Investigation

Submodularity is a set-functional property, intuitively understood as a diminishing returns property. One of several alternative means of proving the property is to show that, for set function $\mathrm{M}(\cdot)$, the following inequality holds for $Y_1, Y_2 \subseteq \mathcal{Y}$ (Nemhauser et al., 1978):

$$\mathrm{M}(Y_1) + \mathrm{M}(Y_2) \geq \mathrm{M}(Y_1 \cup Y_2) + \mathrm{M}(Y_1 \cap Y_2)$$

In this appendix, the submodular properties (or lack thereof) of the conditional entropy and VoI metrics are investigated. The basic modeling assumptions of Section 3.3 are relaxed in this appendix, except where otherwise indicated for specific applications. Furthermore, in this section, sensing costs are ignored in the definitions of sensing objectives. This is done for simplicity of the resulting investigations, such that the properties of these cost functions need not be accounted for. Note that, through an arbitrary definition of a cost function, any objective might be forced to be submodular or otherwise.

B.1. Conditional Entropy

Substituting the reduction in entropy objective of Equation 4-41 (ignoring costs) into the inequality gives:

$$[\mathrm{H}(F) - \mathrm{H}(F|Y_1)] + [\mathrm{H}(F) - \mathrm{H}(F|Y_2)] \geq [\mathrm{H}(F) - \mathrm{H}(F|Y_1 \cup Y_2)] + [\mathrm{H}(F) - \mathrm{H}(F|Y_1 \cap Y_2)]$$

$$\mathrm{H}(F|Y_1 \cap Y_2) + \mathrm{H}(F|Y_1 \cup Y_2) \geq \mathrm{H}(F|Y_1) + \mathrm{H}(F|Y_2)$$

Using the relationship of the conditional and joint entropy, e.g., as $\mathrm{H}(F|Y_1) = \mathrm{H}(F \cup Y_1) - \mathrm{H}(Y_1)$, gives:

$$\begin{aligned}\mathrm{H}(F \cup \{Y_1 \cap Y_2\}) - \mathrm{H}(Y_1 \cap Y_2) + \mathrm{H}(F \cup Y_1 \cup Y_2) - \mathrm{H}(Y_1 \cup Y_2) \\ \geq \mathrm{H}(F \cup Y_1) - \mathrm{H}(Y_1) + \mathrm{H}(F \cup Y_2) - \mathrm{H}(Y_2)\end{aligned}$$

and again using this relationship as $\mathrm{H}(Y_1 \cup Y_2) = \mathrm{H}(Y_2|Y_1) + H(Y_1)$ gives:

$$\begin{aligned}\mathrm{H}(F \cup \{Y_1 \cap Y_2\}) - \mathrm{H}(Y_1 \cap Y_2) + \mathrm{H}(F \cup Y_1 \cup Y_2) - \mathrm{H}(Y_2|Y_1) - H(Y_1) \\ \geq \mathrm{H}(F \cup Y_1) - \mathrm{H}(Y_1) + \mathrm{H}(F \cup Y_2) - \mathrm{H}(Y_2)\end{aligned}$$

which simplifies to:

$$\begin{aligned}\mathrm{H}(F \cup \{Y_1 \cap Y_2\}) + \mathrm{H}(F \cup Y_1 \cup Y_2) + \mathrm{H}(Y_2) \\ \geq \mathrm{H}(F \cup Y_1) + \mathrm{H}(F \cup Y_2) + \mathrm{H}(Y_2|Y_1) + \mathrm{H}(Y_1 \cap Y_2)\end{aligned}$$

It can be shown that this inequality holds in several specific cases. For example, in a trivial case, assume that F, Y_1, and Y_2 are independent. The inequality then becomes:

$$2\mathrm{H}(F) + \mathrm{H}(Y_1 \cap Y_2) + \mathrm{H}(Y_1) + 2\mathrm{H}(Y_2) \geq 2\mathrm{H}(F) + \mathrm{H}(Y_1) + 2\mathrm{H}(Y_2) + \mathrm{H}(Y_1 \cap Y_2)$$

Both sides are equal, so the inequality holds as equality, proving submodularity for independent measurements. However, as this means the measurements are uninformative about the random field, this fact is of little use for supporting optimal sensor placement.

In a more interesting case, assume that Y_1 and Y_2 are subsets of F. Therefore the inequality becomes:

$$2\mathrm{H}(F) + \mathrm{H}(Y_2) \geq 2\mathrm{H}(F) + \mathrm{H}(Y_2|Y_1) + \mathrm{H}(Y_1 \cap Y_2)$$

$$\mathrm{H}(Y_2) \geq \mathrm{H}(Y_2|Y_1) + \mathrm{H}(Y_1 \cap Y_2)$$

Note that $\mathrm{H}(Y_2)$ can be expressed as $\mathrm{H}(Y_2 \backslash \{Y_1 \cap Y_2\}|Y_1 \cap Y_2) + \mathrm{H}(Y_1 \cap Y_2)$. Therefore the inequality becomes:

$$\mathrm{H}(Y_2 \backslash \{Y_1 \cap Y_2\}|Y_1 \cap Y_2) \geq \mathrm{H}(Y_2|Y_1)$$

Also note that $\mathrm{H}(Y_2|Y_1) = \mathrm{H}(Y_2 \backslash \{Y_1 \cap Y_2\}|Y_1 \cup \{Y_1 \cap Y_2\}) + \mathrm{H}(Y_1 \cap Y_2|Y_1) = \mathrm{H}(Y_2 \backslash \{Y_1 \cap Y_2\}|Y_1) + \mathrm{H}(Y_1 \cap Y_2|Y_1) = \mathrm{H}(Y_2 \backslash \{Y_1 \cap Y_2\}|Y_1)$ since $\mathrm{H}(Y_1 \cap Y_2|Y_1) = 0$. This changes the inequality to:

$$\mathrm{H}(Y_2 \backslash \{Y_1 \cap Y_2\}|Y_1 \cap Y_2) \geq \mathrm{H}(Y_2 \backslash \{Y_1 \cap Y_2\}|Y_1)$$

Since $Y_1 \cap Y_2 \subseteq Y_1$, the inequality holds, as conditioning on a larger set of variables always reduces (or does not increase) the entropy. Thus, the property of submodularity holds when observations Y are direct measurements of random variables in F.

More generally, assume that non-overlapping subsets of measurements are conditionally independent given F. First, using the relationship between the joint entropy and conditional entropy allows the inequality to be expressed as:

$$\begin{aligned}\mathrm{H}(Y_1 \cap Y_2|F) + \mathrm{H}(F) + \mathrm{H}(Y_1 \cup Y_2|F) + \mathrm{H}(F) + \mathrm{H}(Y_2) \\ \geq \mathrm{H}(Y_1|F) + \mathrm{H}(F) + \mathrm{H}(Y_2|F) + \mathrm{H}(F) + \mathrm{H}(Y_2|Y_1) + \mathrm{H}(Y_1 \cap Y_2)\end{aligned}$$

which simplifies to:

$$\mathrm{H}(Y_1 \cap Y_2|F) + \mathrm{H}(Y_1 \cup Y_2|F) + \mathrm{H}(Y_2) \geq \mathrm{H}(Y_1|F) + \mathrm{H}(Y_2|F) + \mathrm{H}(Y_2|Y_1) + \mathrm{H}(Y_1 \cap Y_2)$$

Due to the assumption of independence between non-overlapping subsets, $\mathrm{H}(Y_1 \cup Y_2|F) = \mathrm{H}(Y_1 \backslash \{Y_1 \cap Y_2\}|F) + \mathrm{H}(Y_2 \backslash \{Y_1 \cap Y_2\}|F) + \mathrm{H}(Y_1 \cap Y_2|F)$, and so:

$$\begin{aligned}\mathrm{H}(Y_1 \cap Y_2|F) + \mathrm{H}(Y_1 \backslash \{Y_1 \cap Y_2\}|F) + \mathrm{H}(Y_2 \backslash \{Y_1 \cap Y_2\}|F) + \mathrm{H}(Y_1 \cap Y_2|F) + \mathrm{H}(Y_2) \\ \geq \mathrm{H}(Y_1|F) + \mathrm{H}(Y_2|F) + \mathrm{H}(Y_2|Y_1) + \mathrm{H}(Y_1 \cap Y_2)\end{aligned}$$

Note that $\mathrm{H}(Y_1|F) = \mathrm{H}(Y_1 \backslash \{Y_1 \cap Y_2\}|F) + \mathrm{H}(Y_1 \cap Y_2|F)$ and $\mathrm{H}(Y_2|F) = \mathrm{H}(Y_2 \backslash \{Y_1 \cap Y_2\}|F) + \mathrm{H}(Y_1 \cap Y_2|F)$ under this assumption as well, and so:

$$\begin{aligned} 2\mathrm{H}(Y_1 \cap Y_2|F) &+ \mathrm{H}(Y_1 \backslash \{Y_1 \cap Y_2\}|F) + \mathrm{H}(Y_2 \backslash \{Y_1 \cap Y_2\}|F) + \mathrm{H}(Y_2) \\ &\geq 2\mathrm{H}(Y_1 \cap Y_2|F) + \mathrm{H}(Y_1 \backslash \{Y_1 \cap Y_2\}|F) + \mathrm{H}(Y_2 \backslash \{Y_1 \cap Y_2\}|F) + \mathrm{H}(Y_2|Y_1) \\ &+ \mathrm{H}(Y_1 \cap Y_2) \end{aligned}$$

which simplifies to:

$$\mathrm{H}(Y_2) \geq \mathrm{H}(Y_2|Y_1) + \mathrm{H}(Y_1 \cap Y_2)$$

As shown above, this inequality holds true in general. Therefore, submodularity holds when measurements are independent conditional to the random field values. This represents the case where observations Y of F have uncorrelated error.

As promising as these observations are, there are situations where the conditional entropy metric is not submodular. Consider, in contrast to the assumption above, that distinct sets of measurements are marginally independent, but conditionally dependent given the random field values. Start with the following inequality, derived from the above representation using the relationship between joint and conditional entropy:

$$\mathrm{H}(Y_1 \cap Y_2|F) + \mathrm{H}(Y_1 \cup Y_2|F) + \mathrm{H}(Y_2) \geq \mathrm{H}(Y_1|F) + \mathrm{H}(Y_2|F) + \mathrm{H}(Y_2|Y_1) + \mathrm{H}(Y_1 \cap Y_2)$$

Marginal independence between the observation sets indicates that $\mathrm{H}(Y_2|Y_1) = \mathrm{H}(Y_2)$, and so the inequality becomes:

$$\mathrm{H}(Y_1 \cap Y_2|F) + \mathrm{H}(Y_1 \cup Y_2|F) \geq \mathrm{H}(Y_1|F) + \mathrm{H}(Y_2|F) + \mathrm{H}(Y_1 \cap Y_2)$$

Note, however, that in general:

$$\mathrm{H}(Y_1 \cap Y_2|F) \leq \mathrm{H}(Y_1 \cap Y_2)$$

Furthermore, given that the observations are conditionally dependent given F:

$$\mathrm{H}(Y_1 \cup Y_2|F) < \mathrm{H}(Y_1|F) + \mathrm{H}(Y_2|F)$$

Combining these two relationships, it is clear that the inequality related to submodularity above cannot hold. Therefore, the conditional entropy metric is not submodular under this case.

As a demonstrative counterexample, consider single discrete random variable F that takes value 0 with probability of 25%, value 1 with probability 50%, and value 2 with probability 25%. Measurements Y_1 and Y_2 are independent Bernoulli variables, taking value 1 with 50% probability and value 0 with 50% probability. These can be understood as observations of F by assuming that $Y_1 + Y_2 = F$. Therefore, although marginally independent, these observations are conditionally dependent given the value of F. Using the Shannon entropy, the lack of

submodularity can be verified. Noting that $\mathrm{H}(\{Y_1\}) = \mathrm{H}(\{Y_2\}) = 1$, $\mathrm{H}(\{F, Y_1, Y_2\}) = 2$, $\mathrm{H}(\{F\}) = 1.5$, $\mathrm{H}(\{Y_1, Y_2\}|\{F\}) = 0.5$, and $\mathrm{H}(\{F, Y_1\}) = \mathrm{H}(\{F, Y_2\}) = 2$:

$$\mathrm{H}(\{F\}) + \mathrm{H}(\{F, Y_1, Y_2\}) + \mathrm{H}(\{Y_2\}) = 4.5$$

$$\mathrm{H}(\{F, Y_1\}) + H(\{F, Y_2\}) + \mathrm{H}(\{Y_2\}|\{Y_1\}) = 5$$

which contradicts the submodularity inequality.

The use of a sum of marginal entropies formulation will have the same general submodularity properties as the joint entropy, since the inequalities must hold for each of the marginal entropies to ensure that they will hold for the sum.

B.2. Mutual Information

Mutual information between observed and unobserved measurements is in general submodular, as shown by Krause (2008). Using the mutual information, as defined in Equation 4-42, the submodularity inequality is expressed as:

$$\begin{aligned}[\mathrm{H}(\mathcal{Y}\backslash Y_1) - \mathrm{H}(\mathcal{Y}\backslash Y_1|Y_1)] + [\mathrm{H}(\mathcal{Y}\backslash Y_2) - \mathrm{H}(\mathcal{Y}\backslash Y_2|Y_2)] \\ \geq [\mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cup Y_2\}) - \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cup Y_2\}|\{Y_1 \cup Y_2\})] \\ + [\mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cap Y_2\}) - \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cap Y_2\}|\{Y_1 \cap Y_2\})]\end{aligned}$$

which rearranges to:

$$\begin{aligned}\mathrm{H}(\mathcal{Y}\backslash Y_1) + \mathrm{H}(\mathcal{Y}\backslash Y_2) + \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cup Y_2\}|\{Y_1 \cup Y_2\}) + \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cap Y_2\}|\{Y_1 \cap Y_2\}) \\ \geq \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cup Y_2\}) + \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cap Y_2\}) + \mathrm{H}(\mathcal{Y}\backslash Y_1|Y_1) + \mathrm{H}(\mathcal{Y}\backslash Y_2|Y_2)\end{aligned}$$

Using the relationship of conditional and joint entropy, e.g., that $\mathrm{H}(\mathcal{Y}\backslash Y_1|Y_1) = \mathrm{H}(\{\mathcal{Y}\backslash Y_1\} \cup Y_1) - \mathrm{H}(Y_1) = \mathrm{H}(\mathcal{Y}) - \mathrm{H}(Y_1)$, the inequality is expressible as:

$$\begin{aligned}\mathrm{H}(\mathcal{Y}\backslash Y_1) + \mathrm{H}(\mathcal{Y}\backslash Y_2) + \mathrm{H}(\mathcal{Y}) - \mathrm{H}(Y_1 \cup Y_2) + \mathrm{H}(\mathcal{Y}) - \mathrm{H}(Y_1 \cap Y_2) \\ \geq \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cup Y_2\}) + \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cap Y_2\}) + \mathrm{H}(\mathcal{Y}) - \mathrm{H}(Y_1) + \mathrm{H}(\mathcal{Y}) - \mathrm{H}(Y_2)\end{aligned}$$

which simplifies to:

$$\begin{aligned}\mathrm{H}(\mathcal{Y}\backslash Y_1) + \mathrm{H}(\mathcal{Y}\backslash Y_2) + \mathrm{H}(Y_1) + \mathrm{H}(Y_2) \\ \geq \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cup Y_2\}) + \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cap Y_2\}) + \mathrm{H}(Y_1 \cup Y_2) + \mathrm{H}(Y_1 \cap Y_2)\end{aligned}$$

Using the definition of the joint entropy, e.g., $\mathrm{H}(Y_1) = \mathrm{H}(Y_1\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2) + \mathrm{H}(Y_1 \cap Y_2)$ and $\mathrm{H}(Y_1 \cup Y_2) = \mathrm{H}(Y_1 \cup Y_2\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2) + \mathrm{H}(Y_1 \cap Y_2)$, this can be expressed as:

$$\begin{aligned}\mathrm{H}(\mathcal{Y}\backslash Y_1) + \mathrm{H}(\mathcal{Y}\backslash Y_2) + \mathrm{H}(Y_1\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2) + \mathrm{H}(Y_2\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2) + 2\mathrm{H}(Y_1 \cap Y_2) \\ \geq \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cup Y_2\}) + \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cap Y_2\}) + \mathrm{H}(Y_1 \cup Y_2\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2) \\ + 2\mathrm{H}(Y_1 \cap Y_2)\end{aligned}$$

which simplifies to:

$$\begin{aligned}\mathrm{H}(\mathcal{Y}\backslash Y_1) + \mathrm{H}(\mathcal{Y}\backslash Y_2) + \mathrm{H}(Y_1\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2) + \mathrm{H}(Y_2\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2) \\ \geq \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cup Y_2\}) + \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cap Y_2\}) + \mathrm{H}(Y_1 \cup Y_2\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2)\end{aligned}$$

And using this definition again to express $\mathrm{H}(Y_1 \cup Y_2\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2)$ as $\mathrm{H}(Y_1\backslash\{Y_1 \cap Y_2\}|\{Y_2\backslash\{Y_1 \cap Y_2\}\} \cup \{Y_1 \cap Y_2\}) + \mathrm{H}(Y_2\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2) = \mathrm{H}(Y_1\backslash\{Y_1 \cap Y_2\}|Y_2) + \mathrm{H}(Y_2\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2)$:

$$\begin{aligned}\mathrm{H}(\mathcal{Y}\backslash Y_1) + \mathrm{H}(\mathcal{Y}\backslash Y_2) + \mathrm{H}(Y_1\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2) + \mathrm{H}(Y_2\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2) \\ \geq \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cup Y_2\}) + \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cap Y_2\}) + \mathrm{H}(Y_1\backslash\{Y_1 \cap Y_2\}|Y_2) \\ + \mathrm{H}(Y_2\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2)\end{aligned}$$

which simplifies to:

$$\begin{aligned}\mathrm{H}(\mathcal{Y}\backslash Y_1) + \mathrm{H}(\mathcal{Y}\backslash Y_2) + \mathrm{H}(Y_1\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2) \\ \geq \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cup Y_2\}) + \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cap Y_2\}) + \mathrm{H}(Y_1\backslash\{Y_1 \cap Y_2\}|Y_2)\end{aligned}$$

Note that, in general, $\mathrm{H}(Y_1\backslash\{Y_1 \cap Y_2\}|Y_1 \cap Y_2) \geq \mathrm{H}(Y_1\backslash\{Y_1 \cap Y_2\}|Y_2)$ since $Y_1 \cap Y_2 \subseteq Y_2$. Therefore we can subtract the respective terms from both sides and preserve the inequality:

$$\mathrm{H}(\mathcal{Y}\backslash Y_1) + \mathrm{H}(\mathcal{Y}\backslash Y_2) \geq \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cup Y_2\}) + \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cap Y_2\})$$

Expanding $\mathrm{H}(\mathcal{Y}\backslash Y_1)$ and $\mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cap Y_2\})$ using the definition for joint entropy:

$$\begin{aligned}\mathrm{H}(Y_2\backslash\{Y_1 \cap Y_2\}|\mathcal{Y}\backslash\{Y_1 \cup Y_2\}) + \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cup Y_2\}) + \mathrm{H}(\mathcal{Y}\backslash Y_2) \\ \geq \mathrm{H}(\mathcal{Y}\backslash\{Y_1 \cup Y_2\}) + \mathrm{H}(Y_2\backslash\{Y_1 \cap Y_2\}|\mathcal{Y}\backslash Y_2) + \mathrm{H}(\mathcal{Y}\backslash Y_2)\end{aligned}$$

and canceling:

$$\mathrm{H}(Y_2\backslash\{Y_1 \cap Y_2\}|\mathcal{Y}\backslash\{Y_1 \cup Y_2\}) \geq \mathrm{H}(Y_2\backslash\{Y_1 \cap Y_2\}|\mathcal{Y}\backslash Y_2)$$

This inequality holds in general since $\mathcal{Y}\backslash\{Y_1 \cup Y_2\} \subseteq \mathcal{Y}\backslash Y_2$.

B.3. Value of Information

Using the VoI metric, the submodularity inequality is:

$$\mathrm{VoI}(Y_1) + \mathrm{VoI}(Y_2) \geq \mathrm{VoI}(Y_1 \cup Y_2) + \mathrm{VoI}(Y_1 \cap Y_2)$$

Substituting the definition of the VoI from Equation 4-6 and simplifying:

$$\mathbb{E}\mathrm{L}(Y_1) + \mathbb{E}\mathrm{L}(Y_2) \leq \mathbb{E}\mathrm{L}(Y_1 \cup Y_2) + \mathbb{E}\mathrm{L}(Y_1 \cap Y_2)$$

This inequality does not hold in general, except in certain special cases. For example, consider that the measurement sets are redundant, i.e., that $\mathbb{E}\mathrm{L}(Y_1 \cup Y_2) = \mathbb{E}\mathrm{L}(Y_1) = \mathbb{E}\mathrm{L}(Y_2)$. In that case, the inequality simplifies to:

$$\mathbb{E}\mathrm{L}(Y_1) \leq \mathbb{E}\mathrm{L}(Y_1 \cap Y_2)$$

which holds as a consequence of the "information never hurts" principle. However, as this must hold true for any measurement sets, including the empty set, this would imply that $\mathbb{E}\mathrm{L}(Y) = \mathbb{E}\mathrm{L}(\emptyset)$ for any $Y \in \mathcal{Y}$, and thus all measurements would have to have no value.

As a more interesting special case, consider an independent system as in Figure 3-3, where random variables affecting each component, and therefore measurements of these, are independent. Furthermore, consider that this system has a decomposable loss function as in Equation 3-1. First, the submodularity inequality can be expressed following the definition of the expected loss as in Equations 4-1 and 4-5 as:

$$\begin{aligned}\mathbb{E}_{Y_1} \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}_1} \mathrm{L}(\mathbf{f},\mathbf{a}) + \mathbb{E}_{Y_2} \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}_2} \mathrm{L}(\mathbf{f},\mathbf{a}) \\ \leq \mathbb{E}_{Y_1\cup Y_2} \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}_1,\mathbf{y}_2} \mathrm{L}(\mathbf{f},\mathbf{a}) + \mathbb{E}_{Y_1\cap Y_2} \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}_{1,2}} \mathrm{L}(\mathbf{f},\mathbf{a})\end{aligned}$$

where $\mathbf{y}_{1,2}$ denotes the measures of $Y_1 \cap Y_2$. Substituting in the decomposable loss function:

$$\begin{aligned}\mathbb{E}_{Y_1} \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}_1} \sum_{i=1}^{n} \mathrm{L}_i(\mathbf{f}_i,\mathbf{a}_i) + \mathbb{E}_{Y_2} \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}_2} \sum_{i=1}^{n} \mathrm{L}_i(\mathbf{f}_i,\mathbf{a}_i) \\ \leq \mathbb{E}_{Y_1\cup Y_2} \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}_1,\mathbf{y}_2} \sum_{i=1}^{n} \mathrm{L}_i(\mathbf{f}_i,\mathbf{a}_i) + \mathbb{E}_{Y_1\cap Y_2} \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}_{1,2}} \sum_{i=1}^{n} \mathrm{L}_i(\mathbf{f}_i,\mathbf{a}_i)\end{aligned}$$

Using the relationship of Equation 5-4:

$$\begin{aligned}\sum_{i=1}^{n} \mathbb{E}_{Y_1} \min_{\mathbf{a}_i\in\mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_1} \mathrm{L}_i(\mathbf{f}_i,\mathbf{a}_i) + \sum_{i=1}^{n} \mathbb{E}_{Y_2} \min_{\mathbf{a}_i\in\mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_2} \mathrm{L}_i(\mathbf{f}_i,\mathbf{a}_i) \\ \leq \sum_{i=1}^{n} \mathbb{E}_{Y_1\cup Y_2} \min_{\mathbf{a}_i\in\mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_1,\mathbf{y}_2} \mathrm{L}_i(\mathbf{f}_i,\mathbf{a}_i) + \sum_{i=1}^{n} \mathbb{E}_{Y_1\cap Y_2} \min_{\mathbf{a}_i\in\mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_{1,2}} \mathrm{L}_i(\mathbf{f}_i,\mathbf{a}_i)\end{aligned}$$

Now, to account for independence of the measurements across components, introduce the notation that $[\![Y]\!]$ refers to the set of component indices such that measurements in the set Y relate to these components. Note that for $i \notin [\![Y]\!]$ the random variables affecting the component are independent of the measures in Y, and therefore:

$$\mathbb{E}_{Y} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) = \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i)$$

Using this notation, the summations in the inequality are split across measurement sets as:

$$\begin{aligned}
&\sum\nolimits_{i \in [\![Y_1]\!]} \mathbb{E}_{Y_1} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_1} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) + \sum\nolimits_{i \in [\![\mathcal{Y} \setminus Y_1]\!]} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) \\
&\quad + \sum\nolimits_{i \in [\![Y_2]\!]} \mathbb{E}_{Y_2} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_2} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) + \sum\nolimits_{i \in [\![\mathcal{Y} \setminus Y_2]\!]} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) \\
&\quad \leq \sum\nolimits_{i \in [\![Y_1 \cup Y_2]\!]} \mathbb{E}_{Y_1 \cup Y_2} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_1, \mathbf{y}_2} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) + \sum\nolimits_{i \in [\![\mathcal{Y} \setminus Y_1 \cup Y_2]\!]} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) \\
&\quad + \sum\nolimits_{i \in [\![Y_1 \cap Y_2]\!]} \mathbb{E}_{Y_1 \cap Y_2} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_{1,2}} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) + \sum\nolimits_{i \in [\![\mathcal{Y} \setminus Y_1 \cap Y_2]\!]} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i)
\end{aligned}$$

Noting that both the left-hand and right-hand sides have the same number of instances of prior expected loss terms, these can be canceled:

$$\begin{aligned}
&\sum\nolimits_{i \in [\![Y_1]\!]} \mathbb{E}_{Y_1} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_1} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) + \sum\nolimits_{i \in [\![Y_2]\!]} \mathbb{E}_{Y_2} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_2} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) \\
&\quad \leq \sum\nolimits_{i \in [\![Y_1 \cup Y_2]\!]} \mathbb{E}_{Y_1 \cup Y_2} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_1, \mathbf{y}_2} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) \\
&\quad + \sum\nolimits_{i \in [\![Y_1 \cap Y_2]\!]} \mathbb{E}_{Y_1 \cap Y_2} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_{1,2}} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i)
\end{aligned}$$

The summation can be split, e.g., between $i \in [\![Y_2 \setminus \{Y_1 \cap Y_2\}]\!]$ and $i \in [\![Y_1 \cap Y_2]\!]$, such that:

$$\begin{aligned}
&\sum\nolimits_{i \in [\![Y_1]\!]} \mathbb{E}_{Y_1} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_1} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) + \sum\nolimits_{i \in [\![Y_2 \setminus \{Y_1 \cap Y_2\}]\!]} \mathbb{E}_{Y_2} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_2} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) \\
&\quad + \sum\nolimits_{i \in [\![Y_1 \cap Y_2]\!]} \mathbb{E}_{Y_2} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_2} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) \\
&\quad \leq \sum\nolimits_{i \in [\![Y_1]\!]} \mathbb{E}_{Y_1 \cup Y_2} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_1, \mathbf{y}_2} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) \\
&\quad + \sum\nolimits_{i \in [\![Y_2 \setminus \{Y_1 \cap Y_2\}]\!]} \mathbb{E}_{Y_1 \cup Y_2} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_1, \mathbf{y}_2} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i) \\
&\quad + \sum\nolimits_{i \in [\![Y_1 \cap Y_2]\!]} \mathbb{E}_{Y_1 \cap Y_2} \min_{\mathbf{a}_i \in \mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_{1,2}} \mathrm{L}_i(\mathbf{f}_i, \mathbf{a}_i)
\end{aligned}$$

Finally, noting that for $i \in [\![Y_1]\!]$, conditioning on Y_2 has no effect, and vice-versa for $i \in [\![Y_2]\!]$, this can be expressed as:

$$\sum_{i\in[\![Y_1]\!]} \mathbb{E}_{Y_1} \min_{\mathbf{a}_i\in\mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_1} \mathrm{L}_i(\mathbf{f}_i,\mathbf{a}_i) + \sum_{i\in[\![Y_2\setminus\{Y_1\cap Y_2\}]\!]} \mathbb{E}_{Y_2} \min_{\mathbf{a}_i\in\mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_2} \mathrm{L}_i(\mathbf{f}_i,\mathbf{a}_i)$$
$$+ \sum_{i\in[\![Y_1\cap Y_2]\!]} \mathbb{E}_{Y_1\cap Y_2} \min_{\mathbf{a}_i\in\mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_{1,2}} \mathrm{L}_i(\mathbf{f}_i,\mathbf{a}_i)$$
$$\leq \sum_{i\in[\![Y_1]\!]} \mathbb{E}_{Y_1} \min_{\mathbf{a}_i\in\mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_1} \mathrm{L}_i(\mathbf{f}_i,\mathbf{a}_i) + \sum_{i\in[\![Y_2\setminus\{Y_1\cap Y_2\}]\!]} \mathbb{E}_{Y_2} \min_{\mathbf{a}_i\in\mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_2} \mathrm{L}_i(\mathbf{f}_i,\mathbf{a}_i)$$
$$+ \sum_{i\in[\![Y_1\cap Y_2]\!]} \mathbb{E}_{Y_1\cap Y_2} \min_{\mathbf{a}_i\in\mathcal{A}_i} \mathbb{E}_{F_i|\mathbf{y}_{1,2}} \mathrm{L}_i(\mathbf{f}_i,\mathbf{a}_i)$$

where, both sides being equal, the inequality holds, and therefore so does the submodularity of VoI in this particular case.

Another case of interest is that of the VoI defined for Gaussian prediction error. Using the relationship of Equation 5-11, the submodularity inequality is:

$$\mathrm{tr}\left(\boldsymbol{\Sigma}_{F|Y_1}\right) + \mathrm{tr}\left(\boldsymbol{\Sigma}_{F|Y_2}\right) \leq \mathrm{tr}\left(\boldsymbol{\Sigma}_{F|Y_1\cup Y_2}\right) + \mathrm{tr}\left(\boldsymbol{\Sigma}_{F|Y_1\cap Y_2}\right)$$

Using the conditional covariance formula of Equation 3-10, and noting that this conditioning can be performed sequentially for different sets of observations, the inequality becomes:

$$\mathrm{tr}\left(\boldsymbol{\Sigma}_{F|Y_1\cap Y_2} - \boldsymbol{\Sigma}_1\right) + \mathrm{tr}\left(\boldsymbol{\Sigma}_{F|Y_1\cap Y_2} - \boldsymbol{\Sigma}_{F|Y_1\cap Y_2}\boldsymbol{\Sigma}_2\boldsymbol{\Sigma}_{F|Y_1\cap Y_2}\right)$$
$$\leq \mathrm{tr}\left(\boldsymbol{\Sigma}_{F|Y_1\cap Y_2} - \boldsymbol{\Sigma}_1 - \left(\boldsymbol{\Sigma}_{F|Y_1\cap Y_2} - \boldsymbol{\Sigma}_1\right)\boldsymbol{\Sigma}_2\left(\boldsymbol{\Sigma}_{F|Y_1\cap Y_2} - \boldsymbol{\Sigma}_1\right)\right) + \mathrm{tr}\left(\boldsymbol{\Sigma}_{F|Y_1\cap Y_2}\right)$$

where:

$$\boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_{F|Y_1\cap Y_2}\boldsymbol{\Omega}^{\mathrm{T}}_{Y_1\setminus\{Y_1\cap Y_2\}}\boldsymbol{\Sigma}_{Y_1\setminus\{Y_1\cap Y_2\}}{}^{-1}\boldsymbol{\Omega}_{Y_1\setminus\{Y_1\cap Y_2\}}\boldsymbol{\Sigma}_{F|Y_1\cap Y_2}$$

and:

$$\boldsymbol{\Sigma}_2 = \boldsymbol{\Omega}^{\mathrm{T}}_{Y_2\setminus\{Y_1\cap Y_2\}}\boldsymbol{\Sigma}_{Y_2\setminus\{Y_1\cap Y_2\}}{}^{-1}\boldsymbol{\Omega}_{Y_2\setminus\{Y_1\cap Y_2\}}$$

The inequality then simplifies to:

$$\mathrm{tr}(\boldsymbol{\Sigma}_1\boldsymbol{\Sigma}_2\boldsymbol{\Sigma}_1) \leq 2\mathrm{tr}\left(\boldsymbol{\Sigma}_1\boldsymbol{\Sigma}_2\boldsymbol{\Sigma}_{F|Y_1\cap Y_2}\right)$$

When this inequality holds, the VoI metric for Gaussian prediction error will be submodular. The reader is referred to Das and Kempe (2008) for additional details, where it is shown that this metric is submodular in the absence of "suppressor variables", or observations that are independent of the random field until another observation is made.

Appendix C: Parameters for San Francisco Bay Area Example

Table C-A: List of infrastructure components.

Component Name	Relative Coordinates [km]*	Component Type	Failure Cost	Closure Cost
Golden Gate Bridge	(0,0)	Steel Suspension	$1200M	$100M
San Francisco Oakland Bay Bridge (East Span)	(13,0)	Concrete Viaduct	$200M	$150M
San Francisco Oakland Bay Bridge (Middle Span)	(11,0)	New Steel Suspension	$1500M	$100M
San Francisco Oakland Bay Bridge (West Span)	(9,-2)	Steel Suspension	$1100M	$100M
Richmond-San Rafael Bridge (East Span)	(3,13)	Steel Cantilever	$250M	$300M
Richmond-San Rafael Bridge (West Span)	(1,14)	Steel Cantilever	$250M	$300M
San Mateo-Hayward Bridge (East Span)	(24,-24)	Steel Trestle	$300M	$150M
San Mateo-Hayward Bridge (West Span)	(19,-26)	Steel Girder	$400M	$150M
Dumbarton Bridge	(31,-35)	Concrete Girder	$200M	$100M
Carquinez Bridge (East Span)	(22,27)	Steel Cantilever	$300M	$150M
Carquinez Bridge (West Span)	(22,27)	New Steel Suspension	$250M	$100M
Benicia-Martinez Bridge (East Span)	(32,25)	Concrete Segmental	$300M	$50M
Benicia-Martinez Bridge (West Span)	(32,25)	Steel Truss	$200M	$100M
Antioch Bridge	(64,22)	Steel Girder	$200M	$100M
Fruitvale Bridge	(20,-4)	Steel Truss Drawbridge	$10M	$5M
Park Street Bridge	(21,-5)	Steel Cantilever Drawbridge	$15M	$5M
High Street Bridge	(22,-6)	Steel Cantilever Drawbridge	$15M	$5M
Bay Farm Island Bridge	(21,-8)	Steel Truss Drawbridge	$10M	$5M
Caldecott Tunnel	(23,6)	Bored Tunnel	$1600M	$400M
Yerba Buena Island Tunnel	(10,-1)	Bored Tunnel	$10M	$1M
Macarthur Tunnel	(0,-2)	Bored Tunnel	$50M	$5M
Waldo Grade Tunnel	(-1,3)	Bored Tunnel	$25M	$5M
Tom Lantos Tunnel	(-4,-26)	Bored Tunnel	$500M	$100M

Broadway Tunnel	(6,-3)	Cut-and-Cover Tunnel	$5M	$1M
Stockton Street Tunnel	(6,-4)	Cut-and-Cover Tunnel	$5M	$1M
Posey and Webster Street Tubes	(16,-3)	Cut-and-Cover Tunnel	$5M	$1M
Northbrae Tunnel	(20,5)	Cut-and-Cover Tunnel	$5M	$1M

*Relative locations indicate distance east and north from mid-span of the Golden Gate Bridge at (37.820°N, 122.480°W).

Table C-B: List of component capacity parameters. Based on models from HAZUS (2013).

Component Type	**Median [g]**	**Coefficient of Variation**
Steel Suspension	0.7	0.6
Concrete Viaduct	1.1	0.6
New Steel Suspension	1.1	0.6
Steel Cantilever	0.75	0.6
Steel Trestle	0.6	0.6
Steel Girder	1.1	0.6
Concrete Girder	1.3	0.6
Concrete Segmental	1.1	0.6
Steel Truss	0.75	0.6
Steel Truss Drawbridge	1.0	0.6
Steel Cantilever Drawbridge	1.0	0.6
Bored Tunnel	0.8	0.6
Cut-and-Cover Tunnel	0.7	0.6

Table C-C: List of component capacity correlations.

correlation coefficients between capacity variables for different infrastructure component types	Steel Suspension	Concrete Viaduct	New Steel Suspension	Steel Cantilever	Steel Trestle	Steel Girder	Concrete Girder	Concrete Segmental	Steel Truss	Steel Truss Drawbridge	Steel Cantilever Drawbridge	Bored Tunnel	Cut-and-Cover Tunnel
Steel Suspension	0.9	0	0.8	0.2	0.2	0.2	0	0	0.1	0.1	0.1	0	0
Concrete Viaduct	0	0.9	0	0	0	0	0.6	0.6	0	0	0	0	0
New Steel Suspension	0.8	0	0.9	0	0	0	0	0	0	0	0	0	0
Steel Cantilever	0.2	0	0	0.9	0.6	0.6	0	0	0.2	0.1	0.2	0	0
Steel Trestle	0.2	0	0	0.6	0.9	0.7	0	0	0.2	0.1	0.1	0	0
Steel Girder	0.2	0	0	0.6	0.7	0.9	0.2	0	0.2	0.1	0.1	0	0
Concrete Girder	0	0.6	0	0	0	0.2	0.9	0.8	0	0	0	0	0
Concrete Segmental	0	0.6	0	0	0	0	0.8	0.9	0	0	0	0	0
Steel Truss	0.1	0	0	0.2	0.2	0.2	0	0	0.9	0.7	0.4	0	0
Steel Truss Drawbridge	0.1	0	0	0.1	0.1	0.1	0	0	0.7	0.9	0.8	0	0
Steel Cantilever Drawbridge	0.1	0	0	0.2	0.1	0.1	0	0	0.4	0.8	0.9	0	0
Bored Tunnel	0	0	0	0	0	0	0	0	0	0	0	0.7	0.3
Cut-and-Cover Tunnel	0	0	0	0	0	0	0	0	0	0	0	0.3	0.7