Investigating the Relationships Between Material Properties and Microstructural Shapes as Quantified by Moment Invariants

Submitted in partial fulfillment of the requirements for

the degree of

Doctor of Philosophy

in

Materials Science and Engineering

Ryan K.S. Harrison

B.S., Materials Science and Engineering, Washington State University
M.S., Materials Science and Engineering, Carnegie Mellon University

Carnegie Mellon University
Pittsburgh, PA

May, 2018

Abstract

The analysis of microstructural shapes is an underutilized tool in the field of materials science. Typical observations of morphology are qualitative, rather than quantitative, which prevents the identification of relationships between shape and the mechanical properties of a material. Recent advances in the fields of computer vision and high-dimensional analysis have made computer-based shape characterization feasible on a variety of materials.

In this work, the relationship between microstructural shapes, and the properties and function of the material as a whole, is explored using moment invariants as global shape descriptors. A different relationship is examined in each of three material systems: how the three-dimensional shapes of cells in the cotyledons of the plant *Arabidopsis Thaliana* can be used to identify cell function; the two-dimensional shapes of additive manufacturing feedstock powder and the ability to distinguish between images of powders from different samples; and the two-dimensional shapes of γ' precipitates and their influence on the creep resistance of single crystal nickel-base superalloys.

In the case of *Arabidopsis Thaliana* cotyledon cells, three-dimensional Zernike and Cartesian moment invariants were used to quantify morphology, and combined with size and orientation information. These feature sets were then analyzed using unsupervised and supervised machine learning methods. Moderate success was found using unsupervised methods, indicating that natural delineations in the data correlate to cell roles to some degree. Using supervised methods, a success rate of 90% was possible, indicating that these features can be used to identify cell function.

The ability of two-dimensional Cartesian moment invariants to distinguish meaningful features in particles of additive manufacturing feedstock was tested by using these features to classify images of feedstock. Ultimately, simple histogram matching methods were unsuccessful, likely because they rely on the most common particles to draw conclusions. A bag-of-words method was used, which uses high-dimensional visualization and clustering techniques to classify individual particles by common features. Histograms of particle clusters are then used to represent each image. This method was far more successful, and a correct classification rate of up to 90% was found, and comparable rates were discovered

using invariants which describe the shapes only broadly. This indicates that moment invariants are an effective measure of the morphologies of these types of particles, and can be used to classify powder shapes, which control many properties which are relevant to the additive manufacturing process.

In the case of the superalloys, it has been shown that the shape distribution of γ' precipitates can be tracked using second order moment invariants. In addition, several low-order moment invariants are shown to correlate to creep resistance in four alloys examined, which supports the idea that the shape of precipitates plays role in determining creep resistance in these alloys.

ACKNOWLEDGEMENTS

Funding for this work was provided by The Air Force Office of Scientific Research through the Multidisciplinary University Research Initiative, #FA9550-12-1-0458.

I would like to thank my advisor, Professor Marc De Graef, without whom this work would not have been possible. His mentorship, guidance, and patience helped me to grow as a scientist and researcher and allowed my vision for this project to determine its course. I would also like to thank the members of my committee: Professor Elizabeth Holm, Professor Anthony Rollett, and Doctor Jeff Simmons.

I owe much gratitude to my predecessors, Dr. Lily Nguyen and Dr. Patrick Callahan, who offered much advice and guidance. Patrick was particularly helpful and patient with my superalloys research. His efforts made that chapter possible.

I would also like to thank my other academic mentors throughout my journey to this point: Professors David Field, Amit Bandyopadhyay, and Susmita Bose of Washington State University, who helped me take my first steps in materials science research and taught me the difference between an edge and a screw dislocation. Also, Chris Cloke, John Spencer, Tom Alexander, and all my other teachers from Wenatchee High School, Pioneer Middle School, and all the way back to Saint Joseph's Elementary.

Thank you to all of my friends that I have made in Pittsburgh, inside the department and out. You all made my time outside of the office (and in) so much fun. I'm going to miss everyone so much. Big thanks to Yanning Zhang who showed a lot of selflessness in helping me after my move to Pittsburgh. I'm not sure I would have gotten to this point without him. Also the folks at St. Margaret who showed me such support over these four years, rain or shine.

My grandparents, to whom my education was so important, and who instilled in me a

curiosity for the world from an early age and who supported me all the way. Thanks to my parents, Mike and Pam, who taught me the importance of education, hard work, and made me the person that I am today. Their support and belief in me never wavered, even when mine did. I owe them everything. Also to the rest of my family, who have given me guidance and advice throughout the years. Finally, I owe so much to Carriann, who has been by my side as we have both worked towards our doctorates. Her endless patience and support have been so meaningful to me in so many ways. I hope I have helped her as much as she has helped me.

List of Publications and Presentations Resulting from this Work

The following papers and presentations related to the work presented in this thesis have been published or submitted:

1. Ryan Harrison and Marc De Graef. *γ' Shape Analysis Using Higher Order Moment Invariants.* Presentation at Materials Science and Technology, Columbus, OH, October, 2015.

2. Ryan Harrison and Marc De Graef. *On the Use of Higher Order Moment Invariants in the Classification of Microstructural Shapes.* Presentation at TMS Meeting and Exhibition, Nashville, TN, February, 2016.

3. Ryan Harrison, Marc De Graef, Patrick Callahan, and Tresa Pollock. *γ' Precipitate Shape and Mechanical Properties in Ni-Base Superalloys.* Poster at 3D Materials Science, St. Charles, IL, July, 2016.

4. Ryan Harrison, Marc De Graef, Patrick Callahan, and Tresa Pollock. *Characterizing γ' Shape Evolution in Nickel-Base Superalloys Using Lower Order Moment Invariants.* Presentation at TMS Meeting and Exhibition, San Diego, CA, February, 2017.

5. Ryan Harrison, Marc De Graef. *Automatic shape-based cell identification in Arabidopsis Thaliana cotyledons using 3D moment invariants.* Presentation at TMS Meeting and Exhibition, Phoenix, AZ. March, 2018.

6. Ryan Harrison, Saransh Singh, Marc De Graef. Orientation similarity of EBSD patterns through image analysis, high-dimensional mapping, and clustering. *Microscopy and Microanalysis.* In preparation, 2018.

7. Ryan Harrison, Marc De Graef. Automatic cell identification in Arabidopsis Thaliana cotyledons using 3D moment invariants, volume, and orientation features. *Acta Biomaterialia.* In preparation, 2018.

8. Ryan Harrison, Brian De Cost, Elizabeth Holm, Marc De Graef. Identifying additive

manufacturing powder images through particle morphologies, as measured by two-dimensional moment invariants. In preparation, 2018.

TABLE OF CONTENTS

Abstract . . . . . **i**

Acknowledgements . . . . . **ii**

List of Publications Resulting from this Work . . . . . **v**

List of Tables . . . . . **x**

List of Figures . . . . . **xi**

I. Introduction 1
- 1.1 Motivation 1
- 1.2 Hypothesis 4
- 1.3 Structure of this document 6

II. Shape and shape descriptors 7
- 2.1 Shape descriptor criteria 8
- 2.2 Two-dimensional shape descriptors 10
  - 2.2.1 Contour-based structural descriptors 11
  - 2.2.2 Contour-based global descriptors 16
  - 2.2.3 Region-based structural descriptors 22
  - 2.2.4 Region-based global descriptors 25
  - 2.2.5 Summary 27
- 2.3 Three-dimensional shape descriptors 28
  - 2.3.1 Feature based methods 28
  - 2.3.2 Graph based methods 30
  - 2.3.3 Geometry based methods 32
  - 2.3.4 Summary 33

III. Moments and moment invariants in shape classification 34
- 3.1 Cartesian moments and invariants 37
  - 3.1.1 Two-dimensional Cartesian moments 38
  - 3.1.2 Cartesian Moment Invariants 39
  - 3.1.3 Three-dimensional Cartesian moments 44
- 3.2 Zernike moments and invariants 45
  - 3.2.1 Two-dimensional Zernike moments 45
  - 3.2.2 Three-dimensional Zernike moments and invariants 48

IV. Cell function identification in *Arabidopsis Thaliana* cotyledon through 3D shape analysis 49
- 4.1 Background 49

4.1.1 Shape analysis in cell and tissue analysis 49
4.1.2 Arabidopsis Thaliana 53
4.2 High dimensional analysis 55
4.2.1 Supervised machine learning 56
4.2.2 Unsupervised machine learning 57
4.3 Materials and methods 60
4.3.1 *Arabidopsis Thaliana* cotyledons 60
4.3.2 Images and image processing 61
4.3.3 Ground truth and scoring 62
4.3.4 Feature sets 63
4.3.5 Unsupervised machine learning 65
4.3.6 Supervised machine learning 66
4.4 Results 66
4.4.1 Mapping and clustering 66
4.4.2 Support vector machine analysis 77
4.5 Discussion 88
4.5.1 Unsupervised analysis 88
4.5.2 Supervised analysis 92
4.5.3 Comparing the approaches 97
4.5.4 Generalized classifiers 99
4.6 Conclusions 99
V. Shape-based classification of metal powders for additive manufacturing 101
5.1 Background 101
5.1.1 Additive manufacturing 101
5.1.2 Particle shape and material properties and defects 103
5.1.3 Powder characterization 104
5.2 Materials and methods 106
5.2.1 Segmentation and feature extraction 108
5.2.2 Data analysis 108
5.3 Results 113
5.3.1 SOMIM and PMIM matching 113
5.3.2 Bag-of-words 117
5.4 Discussion 129
5.4.1 Map matching 129
5.4.2 Bag-of-words 130
5.4.3 Comparing approaches 133
5.4.4 Comparisons and applications 135
5.5 Conclusions and future work 136
VI. γ' precipitate shape and creep resistance in single-crystal nickel-base superalloys 138
6.1 Background 138
6.1.1 Nickel-base single crystal superalloy microstructure and properties 138
6.1.2 Shape descriptors in superalloy analysis 146
6.2 Materials and methods 149
6.2.1 Sample preparation 149
6.2.2 Image processing and segmentation 150

6.2.3 Creep sample preparation 154
6.3 Results 154
6.3.1 Heat treatments and precipitate growth 154
6.3.2 Creep results 162
6.4 Discussion 173
6.4.1 Aging 173
6.4.2 Creep and precipitate morphology 175
6.5 Conclusions and future work 180
6.5.1 Future work 181
VII. Summary and Conclusions 182
7.1 Summary 182
7.2 Conclusions 184

Bibliography **185**

LIST OF TABLES

4.1 Properties of feature sets 64
4.2 Results of unsupervised clustering on both cotyledon data sets. Standard deviation of NMI calculated over five iterations for each feature set for the set of clustering parameters with highest average NMI. 67
4.3 Results of multiple linear correlation analysis performed on variables affecting clustering results. Dependent variable is the mean mutual information for five iterations. r_l is the linear correlation coefficient for the given parameter, while R^2_m is the multiple linear correlation coefficient. ΔR^2_m is the squared semipartial correlation for the given feature. 72
4.4 Summary of support vector machine analysis of 7x and 8x datasets via five-fold cross validation. 78
4.5 Results of multiple linear correlation analysis performed on indicated feature sets for SVM. Dependent variable is the percent success from five-fold cross validation over five iterations. "secord" indicates second order moments (orientation). r_l is the linear correlation coefficient for the given parameter, while R^2_m is the multiple linear correlation coefficient. ΔR^2_m is the squared semipartial correlation for the given feature. 82
4.6 Summary of cross-cotyledon application of models 87

5.1 Parameters of SLS/SLM that affect the final properties and the presence of defects in a final part. Modified from [1]. 102

6.1 Alloys used and compositions of each (weight percent). 149
6.2 Heat treatments for all alloys examined in this study. 150
6.3 Summary of mean shape values and creep results for the examined alloys. 163

LIST OF FIGURES

2.1 Allowed grid points and relative vector labels and directions for two different chain code schemes. From [2]. 11
2.2 Input data, features, and feature index tree from [3] 13
2.3 Apple silhouette and resulting centroid distance shape signature. From [4] 18
2.4 Ape skulls, resulting shape signatures, and resulting hierarchy of similarities. From [5] 19
2.5 Formation of medial axes using grassfire transform 23
2.6 Matching shock graphs despite occlusion or missing regions. Same color shocks represent matches. From [6]. 24
2.7 Example of Reeb graphed based on height along z axis for a torus. From [7]. 30

3.1 Weight function for basis function x^p with increasing values of p. 35
3.2 The relative values of the basis function $x^p y^q$ for the given (p, q). Values shown for $-1 \leq x \leq 1$ and $-1 \leq y \leq 1$. 38
3.3 Reconstruction of the letter "E" using geometric moments. From [8]. . . . 39
3.4 Second order moment invariant map (SOMIM) generated from plotting ω_1 against ω_2 normalized by the values of the cirlce. From [9]. 43
3.5 Radial function (R_{nm}) for first ten Zernike polynomials. 47
3.6 Reconstruction of of a spaceship model using 3D Zernike moments. From [10]. 48

4.1 Examples of 7x series images. (a-c) example images from along the z-axis. (d) shows a image taken along the x-axis with more cylindrical palisade mesophyll visible along the bottom of the cotyledon. 54
4.2 Examples of 8x series images. 54
4.3 ANOVA table and histogram of residuals for the influence of feature sets and algorithm parameters on unsupervised clustering results on 7x data set. 69
4.4 ANOVA table and histogram of residuals for the influence of feature sets and algorithm parameters on unsupervised clustering results on 8x data set. 70
4.5 Residual plots for 7x and 8x respectively after interaction-free multiple linear regression. There is a clear bimodal distribution of points and a negative trend, indicating that some factor not listed is likely at play, or interactions, which have not been included, are having an effect. 72
4.6 Similar cells identified after initial t-SNE and HDBSCAN clustering (left) and tree merging (right) of 7x dataset with CVNS feature set, perplexity of 30, min_clust of 50, and min_samp of 35. Clustering shown has an NMI 0f 0.420. 74

4.7 Generated t-SNE mapping and HDBSCAN clustering on 7x dataset with CVNS features, perplexity of 10, min_clusts of 50, and min_samples of 10. Clustering shown has an NMI 0f 0.346. 74
4.8 Identified clusters based upon t-SNE mapping and HDBSCAN clustering of Figure 4.6. 75
4.9 Identified clusters based upon t-SNE mapping and HDBSCAN clustering of Figure 4.7. 76
4.10 Results of SVM grid search as measured by five-fold cross validation percentage for 7x data set. 79
4.11 ANOVA table and histogram of residuals for the influence of feature sets supervised, SVM classifier on 7x data set. 80
4.12 ANOVA table and histogram of residuals for the influence of feature sets supervised, SVM classifier on 7x data set. 81
4.13 Residual plots for the multiple linear regression analysis of Table 4.5 . . . 82
4.14 Example of confusion matrix, shown here larger to illustrate the position of the categories and axis titles. 83
4.15 Set of confusion matrices for SVM of 7x data set 84
4.16 Set of confusion matrices for SVM of 8x data set 85
4.17 Identified clusters based upon SVM classifier for 7x dataset generated by CVNS feature set. 86
4.18 Comparison of cross validation success between cotyledons in t-SNE mapping and HDBSCAN clustering with indicated feature sets.Values taken from Table 4.2 89
4.19 Comparison of cross validation success between cotyledons in SVM with indicated feature sets.Values taken from Table 4.4 93

5.1 Examples of powder images used in this analysis. Courtesy of the Holm Group, Carnegie Mellon University. Degree of magnification and therefore particle size was not included in this analysis. 107
5.2 Example moment invariant maps with identified shape regions. 110
5.3 PCA versus t-SNE mapping, from 80-dimensional Cartesian moment invariant space to two-dimensions, of the over 50,000 particles analyzed in this chapter. 112
5.4 Master SOMIMs for each class of powder material. 115
5.5 Master PMIMs for each class of powder material. 116
5.6 Confusion matrix for Hellinger distance mapping of test image SOMIMs to overall material SOMIMs 118
5.7 Confusion matrix for Hellinger distance mapping of test image PMIMs to overall material PMIMs 118
5.8 Confusion matrix for Hellinger distance mapping of both SOMIM and PMIMs 119
5.9 t-SNE map. Perplexity of 50. 120
5.10 Example of results from t-SNE mapping from five-dimensional ω-space to two-dimensions. 121
5.11 t-SNE map. Perplexity of 50. 122
5.12 Example of results from t-SNE mapping from eighty-dimensional Cartesian moment invariant-space to two-dimensions. 124
5.13 Confusion matrix for ω invariants, matching image histograms to material histograms. 125

5.14 Confusion matrix for ω invariants, matching image histograms to material histograms, with distance-based voting for the k-NN classifiers. 125
5.15 Confusion matrix for ω invariants, matching image histograms to training image histograms. 126
5.16 Confusion matrix for ω invariants, matching image histograms to training image histograms, with distance-based voting for the k-NN classifier. . . . 126
5.17 Confusion matrix for Cartesian invariants, matching image histograms to material histograms. 127
5.18 Confusion matrix for Cartesian invariants, matching image histograms to material histograms, with distance-based voting for the k-NN classifiers. . 127
5.19 Confusion matrix for Cartesian invariants, matching image histograms to training image histograms. 128
5.20 Confusion matrix for Cartesian invariants, matching image histograms to training image histograms, with distance-based voting for the k-NN classifier. 128
5.21 Mapping t-SNE and HDBSCAN clusters back to a PCA mapping, giving an idea of the location of these clusters. 134

6.1 The $L1_2$ structure, shown for Ni_3Al. From [11] 140
6.2 The growth of γ' precipitates. From [12] 141
6.3 Segmentation examples using *ilastik* software. 152
6.4 Segmentation comparison of EM/MPM algorithm and *ilastik* segmentation method. 153
6.5 Segmented microstructural evolution of γ' precipitates for indicated aging times, in hours. All images are $5.18\mu m$ across. 156
6.6 Second order moment invariant maps for indicated aging times, in hours. . 158
6.7 Box plots of ω_1 for all measured precipitates. Asterisks indicate mean. . . 159
6.8 Box plots of ω_2 for all measured precipitates. Asterisks indicate mean . . . 160
6.9 Trends in aging, with equivalent volume of precipitates plotted on the y-axis, and aging time in hours on the x. Strong linear relationships indicate adherence to LSW growth. Blue lines indicate selection of heat treatment times for creep testing. 161
6.10 Micrographs of superalloy microstructures prior to creep testing. All images have a width of $15\mu m$. 164
6.11 SOMIMs for creep-tested alloys. 165
6.12 PMIMs for creep-tested alloys. 166
6.13 Kernel density estimates for the four examined materials and the indicated properties. Bars above the plot indicate overall mean as a diamond, plus or minus the standard deviation of the data. The square symbols indicate the position of the mean of shapes which have a cuboidal morphology. Vertical dashed lines indicate the median of the data. 167
6.14 Stress versus time in seconds for the alloys examined here. ME1 showed the greatest time to failure and overall creep resistance while CE5 showed the least. 168
6.15 Best fit line for overall mean of the features plotted against the natural logarithm of the minimum creep rate for each alloy. Colors of symbols indicate the alloy, matching those of Figure 6.14. 169

6.16 Best fit line for the median values of the features plotted against the natural logarithm of the minimum creep rate for each alloy. Colors of symbols indicate the alloy, matching those of Figure 6.14. 170
6.17 Best fit line for the mode values, taken from KDEs of Figure 6.13 of the features plotted against the natural logarithm of the minimum creep rate for each alloy. Colors of symbols indicate the alloy, matching those of Figure 6.14. 171
6.18 Post-creep micrographs. Stress axis is normal to the images. Lighter phase is γ'. 172

CHAPTER I

Introduction

1.1 Motivation

The relationships and interdependencies between the processing, structure, and properties of a material comprise what is commonly referred to as the materials science paradigm. *Processing* refers to the conditions under, and methods by which, a material is synthesized or manipulated. It is through processing that influence is exerted on the structure of the material in question. *Structure* describes the physical state of the matter being investigated. *Properties* of the material are quantifications of the way in which the material reacts to or is influenced by the world around it. Knowledge of all parts of the paradigm allows a material to be optimally utilized by identifying a desirable property or set of properties, determining the structure that will lead to these properties, and implementing the proper processing to lead to this structure. Without knowledge of structure, it is impossible to know why or from where these properties originate.

Characterization is the toolbox through which knowledge of structure is built. Over the past seven decades, the desire for greater understanding of the structure/property relationship has driven an explosion in the number of characterization techniques available to the materials scientist, and equivalently, the different ways of quantifying a material. There have been advancements in understanding and measurement of nearly every aspect of a material's structure, including crystallographic orientation, physical structure through improved microscope resolution, topography, composition, atomic organization, and three dimensional spatial information.

While shape characterization has been utilized in materials science in the past, especially

in two-dimensions, shape has remained an infrequently examined facet of microstructure. This may be because shape has been commonly used in a qualitative way, generally preventing any kind of statistical analysis. Over the history of the field, the vast majority of morphological information has been noted by using poorly defined qualitative descriptors such as "acicular", "lamellar", and "circular". This is especially true with regard to microstructural analysis in metallurgy. While these labels may be sufficient in the context of the work in which they appear, especially if morphology is not the primary subject of examination, they stymie any further investigation into shape and shape/property relationships in the material. These qualitative descriptors can even mask subtle variations in shape. When an investigator identifies a subset of the microstructure by this type of label, differences in aspect ratio, complexity, or other less obvious morphological information are often ignored. This is counterproductive to an investigation into a material which seeks to fully understand the nature of the material being investigated.

When shape *is* measured quantitatively, it is primarily through constructed parameters which may have no meaning outside of the current system, or a certain subset of expected morphologies. A typical example of this is the comparison of the perimeter to area ratio of a shape in question to that of an reference shape such as a square or circle. For descriptors such as these, it is possible to generate identical values using shapes which differ dramatically. These values may be sufficient for studying a basic range of shapes, but lose the ability to discriminate between shapes when they deviate from the expected morphologies. At best, these types of quantifiers capture only one aspect of a particle's shape. At worst, they obfuscate information and encourage investigators to focus on and only measure expected morphologies within the material. These types of parameters will be more completely discussed later in Chapter II, along with other methods of measuring and quantifying morphology.

A likely factor behind the lack of detailed and complete shape analysis in materials science is the difficulty in gathering shape information on large numbers of particles. Even in recent publications, many measurements are taken by hand. Manual measurements dramatically reduce the number of measured shapes which can be observed per sample due to the time it takes to perform the measurements. In addition, manual measurement

reduces the variety of shape information that can be captured. Properties such as Feret diameter are measured relatively easily, but even simple values like area or perimeter can become difficult to measure for more irregular shapes.

Despite the many issues with manual measurements, it has remained popular due to the difficulty associated with implementing computer based solutions. The most appealing feature of manual measurement may be that the individual doing the measuring is free to identify the boundaries of regions of interest, in effect segmenting the object by eye. The process of performing an automatic segmentation has remained a difficult problem since the implementation of computer-based graphics in the 60s [13]. However, relatively recent advances in computer science, especially in the fields of computer vision and machine learning have improved the reliability and accuracy of the segmentation process [14–16], and opened the way for computer based shape analysis of images in materials science.

Solving the second problem of manual measurements: limited shape information, relies on using a morphology metric that captures many aspects of an object's shape. Moment invariants are a useful tool in this regard. They can be calculated up to any desired order, each of which contains information regarding finer and finer details of an object's shape [17]. In addition, lower order moment invariants can be used to capture intuitive shape information that allows for easy interpretability. Calculating moment invariants up to a large degree of accuracy for a large number of shapes can be resource intensive. However, with modern equipment, including the capacity for parallelization, the time required can be reduced and these calculations made possible on a relatively large scale.

Shape analysis utilizing a large number of moment invariants, while generally being more accurate, necessarily occurs in high dimensional spaces. These types of spaces are inherently difficult to visualize, and even more difficult to work in. Even relatively simple measurements such as distance become unreliable. However, high dimensional visualization and clustering techniques are becoming more sophisticated with time. These advancements have made the creation of two- and three- dimensional projections of high dimensional spaces more accurate by allowing the structure of the data to survive the process [18, 19]. Combined with improvements in clustering algorithms, these techniques have improved the quality of information gleaned from high dimensional spaces, and allow improvements in

shape analysis through the analysis of many aspects of shape at once.

Moment invariants can also be expanded into three dimensions. Three dimensional materials analysis has recently come to the fore with advances in characterization tools. Techniques such as plasma focused ion beam SEM microscopy, femtosecond laser ablation, and advances in tomography have dramatically increased the size of three-dimensional regions which can be analyzed, and reduced the time necessary to gather information [20,21]. Three dimensional moment invariants can take advantage of these improved techniques while also providing an analogue to their 2D counterparts. A full discussion of shape parameters is given in Chapter II.

With advances in technologies and techniques over the last few decades, particularly in computer science, the way has been cleared for a more advanced and nuanced quantification of shape than ever before. The goal of this work is therefore to present the use of this kind of shape analysis in materials science, both in two and three dimensions, and to illustrate how this knowledge can provide another dimension through which the structure/properties relationship can be viewed and understood.

1.2 Hypothesis

Shapes of subregions in a material are typically the result of complicated interactions between a variety of invisible or difficult to observe factors, such as stress fields or composition gradients. These shapes can therefore serve as a visible indicator of the state of these properties, and also exert their own influence on the material. In order to explore the interactions between shape and properties, a number material systems are examined, with the morphologies of different microstructural features on different scales being measured and compared to material function for each. Measured features and systems include: two dimensional γ' precipitate shapes in experimental single-crystal nickel-base superalloys, and their correlation to creep resistance of the bulk material; three dimensional cell shapes in cotyledons of *Arabidopsis Thaliana*, and the identification of cell function; and powder feedstock morphology in two and three dimensions prior to use, and final bulk material properties for a number of different alloy powders for additive manufacturing.

The hypothesis of the research presented in this document is: *Shapes of microstruc-*

tural subregions in embryonic Arabidopsis Thaliana *leaves, additive manufacturing feedstock powders, and single-crystal nickel-base superalloys, as quantified by Cartesian and Zernike moment invariants, correlate to the overall properties or function of the material as a whole.* This hypothesis will be explored in the following material systems under the given conditions:

1. The three-dimensional shapes of cells in *Arabidopsis Thaliana* cotyledons, quantified through three-dimensional Cartesian and Zernike moment invariants, can be used to identify cell function. This can be achieved through both unsupervised and supervised machine learning methods.

2. Two-dimensional Cartesian moment invariants up to the 12^{th} order quantify relevant morphological information of additive manufacturing feedstock particles. This will be examined through identification of powder images by matching moment invariant distributions to known images, indicating if moment invariants are capable of making meaningful distinctions between the classes of shapes present in these materials.

3. The two-dimensional shapes of γ' precipitates in single crystal nickel-base superalloys, quantified with Cartesian moment invariants up to the 4^{th} order, can be correlated to the creep resistance of the alloy.

Utilizing a variety of material systems serves multiple purposes with regard to proving this hypothesis. The first is that it shows that shape analysis by moment invariants is a useful tool that can be applied across materials systems. Given the diversity of systems, it is likely that any shape/property relationships found here do not share the same underlying mechanisms; shape can influence properties in a number of different ways which must be taken on a case by case basis.

Examining a variety of shapes also says something about the shape descriptors themselves. There is no guarantee that moment invariants are appropriate for every class of shapes, or that they are capable of quantifying relevant shape features for every examined system. However, by working with many material systems the full extent to which moment invariants can quantify relevant shapes can be explored.

1.3 Structure of this document

Chapter II contains an discussion of shape, categories of shape descriptors, and their relative strengths and weaknesses in terms of object identification, as well as examples of each. Chapter III provides information on moments and moment invariants of the type used in this work, with a focus on the underlying mathematics.

Chapters V, IV, and VI detail the specific experiments undertaken to prove the hypothesis. Each chapter contains a brief introduction which provides specific information relevant to the problem at hand, followed by the methodology undertaken for the given experiment. This is followed by results, and a discussion of the implications of what has been found. Each of these chapters is concluded with a summary of the results, their relationship to the hypothesis and a statement of the future work which may be necessary to further explore additional questions. Finally, Chapter VII reiterates the main points of this entire work, how they relate to the hypothesis, and summarizes additional future work which may help further explore and illuminate questions regarding the hypothesis.

CHAPTER II

Shape and shape descriptors

Shape is the form of an object; the relative location of all of its component area or volume elements. According to Kendall [22], shape can also be thought of as what it is not: it is what is left when "location, orientation, and scale" information are removed from an object. Two subsets of Euclidean space have the same shape if they can be rotated, translated and scaled to be identical to each other. At first glance, shape appears to be a simple concept, and human beings are able to identify shapes that are similar or dissimilar from an early age. However, this understanding seems to be more instinctual than analytical and begins to break down when pushed to describe greater level of detail. When asked if shape A is more or less similar to shape B or C, the response is likely to be, "In what way?" This question demonstrates the need for numerical shape descriptors and simultaneously illustrates their usefulness. Not only do they provide a way to quantify shape, and provide an answer to which shapes are more similar than others, they also supply a method to break the morphology of an object into different aspects. By utilizing these descriptors, the relevant facets of shape can be determined and analyzed, and conclusions can be drawn about each descriptor separately from each other.

The most commonly used shape descriptors are simple qualitative descriptions. Moreover, these qualitative descriptions are typically applied to entire collections of shapes. Grains in a material which have different morphologies will be labeled as "plate-like", or precipitates which span a range of aspect ratios will be deemed "circular" without any additional analysis. This qualitative method is quick and allows simple visual identification, but it is prone to error and can falsely label shapes that are dissimilar as identical. Be-

cause it provides only a single value as a measurement, this technique also prevents any dissection of the various aspects of morphology. When shape is not the primary property in question, this analysis can be useful for the purposes of drawing simple conclusions, but it obviously prevents any kind of deeper statistical analysis into the effects that shape can have over the behavior of a material. Descriptors are needed that address the weaknesses of these common labels.

2.1 Shape descriptor criteria

There is no such thing as an ideal shape descriptor, what is clear is that a good parameter should be a similarity measure: it should not only quantify the morphology of an object, but also proscribe similar values to shapes which share similar features. In addition, a useful shape descriptor should have several other properties. The first is that the descriptor should be capable of measuring multiple aspects of a shape, for example, aspect ratio and circularity. By separating different facets of shape, it makes determination of similarity more obvious. Ideally these various aspects correspond to properties of morphology which are, at least at gross levels of examination, clear to human perception. This may help to ensure that the shapes which are perceived as similar to the naked eye can be classed as similar by the similarity measures as well.

To calculate similarity, it is important that the space defined by the quantifiers is a *metric space*. A metric space is an ordered pair (X, d) where X is any set and d is a *metric* where, $d : X \times X \to \mathbb{R}$. The metric must also satisfy the following four conditions as listed in Biasotti *et al.*, where $x, y, z \in X$:

$$d(x, y) \geq 0 \qquad \text{(non-negativity)} \tag{2.1a}$$

$$d(x, y) = 0 \text{ if and only if } x = y \qquad \text{(identity)} \tag{2.1b}$$

$$d(x, y) = d(y, x) \qquad \text{(symmetry)} \tag{2.1c}$$

$$d(x, y) \leq d(x, z) + d(z, y) \qquad \text{(triangle inequality)} \tag{2.1d}$$

These conditions ensure that the space being utilized behaves similarly to Euclidean space, a metric being the generalization of the behavior of Euclidean distance. Equation 2.1a

ensures that any distance measured in this space is positive. Equation 2.1b implies that if two points occupy the same location, they are identical. Equation 2.1c asserts that the distance measure is not a one-way measurement, while equation 2.1d enforces that the shortest distance between two points is a straight line.

Marr and Nishihara [23] identified three criteria by which to measure the usefulness of a shape representation. These are *accessibility*, *scope and uniqueness*, and *stability and sensitivity*. *Accessibility* refers to the ease of extracting a shape descriptor from a given image. This includes computation time as well as limitations of the descriptor or the data, such as resolution. *Scope and uniqueness* indicates the range of shapes which can be accurately represented by the descriptor. For a descriptor designed to measure cuboidal cross-sections, a square or circle would be within its *scope*, but not a triangle. *uniqueness* is the ability of a descriptor to give the same value for the same shape, regardless of its coordinate system. Finally, *stability and sensitivity* represents two competing but important concepts. *Stability* is used to indicate the way a descriptor should give two similar shapes similar values, while *sensitivity* refers to the ability of the descriptor to differentiate subtle differences between similar shapes.

Brady [24] proposed three additional criteria to those of Marr. First, a representation should be information preserving, meaning that no shape information is lost during its computation. It should also be computed locally: on a scale less than that of the shape as a whole. Together these are referred to as *rich local support*. Descriptors which display *smooth extension and subsumption* are those which can be used to create global descriptors (descriptors which capture aspects of the object as a whole.Finally, *propagation* implies that visually identifiable subparts of the shape which are captured locally are also captured globally.

Similar criteria is given in Zhang and Lu [4], who identify as desirable the features of: invariance to transformations of the object, insensitivity to noise or distortion of the object, a hierarchical nature which allows finer or coarser details to be quantified, application independence (the capability to represent many shapes meaningfully), and low computation complexity. Woodham [25], Binford [26], and Mokhtarian *et al.* [27] also conduct analyses on shape descriptors and give ideal properties similar to the other discussed here.

Based upon this information, necessary properties for an effective shape descriptor to be utilized in this work must include:

1. Application to a wide variety, if not all, two or three dimensional shapes
2. Invariance to translation, rotation and scaling
3. Returns similar values for similar shapes while still retaining the ability to discriminate between like objects
4. Has a hierarchical structure which allows for the analysis of broader or more narrow shape information
5. Values which reflect visual shape information
6. Generates a metric space

With these qualities in mind, what follows is an overview of the types of shape descriptors that are commonly used to quantify shape. It is not meant to be an exhaustive list, but instead a representative cross-section of the categories of techniques that are available to use when quantifying shape in a number of different fields. Not all descriptors discussed here necessarily fulfill all of the requirements above.

2.2 Two-dimensional shape descriptors

Two dimensional shape analysis is the most common kind of shape analysis performed, primarily due to the fact that the only input data necessary to perform this kind of analysis is a two-dimensional image. Digital images are easy to store, simple to work with, and reflect how humans see the world: as a two-dimensional projection of a three-dimensional space. The relative ease of capturing 2D images as compared to full three-dimensional representations also contributes to an asymmetry in the types of available data. Two dimensional shape analysis has also garnered attention in the past decade because of its importance in computer vision, which is contributing to current advances in the automation of a number of industries.

In their review of two-dimensional shape descriptors, Zhang and Lu identified two major categories of these quantities: region-based descriptors, which capture information regarding the region contained by the shape, and contour-based descriptors, which only utilize

the boundary of the object. They further categorized these descriptors by identifying either global, or structural quantities. Global indicates that the value measures the entire shape as a whole, while a structural parameter breaks the object into component parts prior to evaluation [4]. In addition to these categories, Loncaric [28] describes two additional distinctions: the first between *scalar transform* techniques which return scalar or vector values, and *space-domain* techniques which return other images; the second distinction is between *information preserving* methods, which allow for complete reconstruction of the generating object, and *non-preserving* techniques, where information is lost and reconstruction is impossible.

2.2.1 Contour-based structural descriptors

Contour-based structural descriptors decompose the boundary of the object into sub-boundaries or *primitives*, which, when taken together, recreate the entire boundary as closely as possible. The shape can then be represented as a string of primitives. A common and popular example is chain code, first introduced by Freeman [2]. With chain code, the boundary is represented by a number of vectors of a fixed length, and only a small number of predefined directions. These vectors, represented by a number, make up a single string of numbers which defines the boundary of the shape. Starting from an arbitrary location, continue to follow the direction of the next vector, which defines an enclosed shape or a line. Figure 2.1

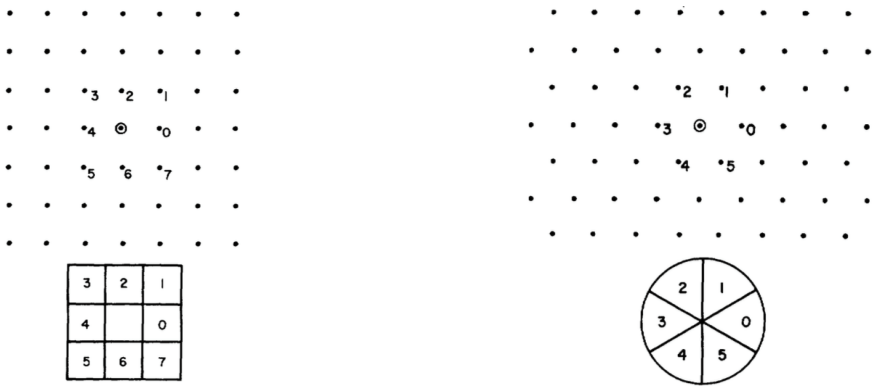

(a) Square grid structure (b) Hexagonal grid structure

Figure 2.1: Allowed grid points and relative vector labels and directions for two different chain code schemes. From [2].

shows two possible vector conventions. The square convention shown in Figure 2.1a uses

eight vectors each corresponding to directions, starting at horizontally to the right, and moving counterclockwise 45° at a time. An alternate convention shown in Figure 2.1b uses 60° steps, which reduces the possible resolution of the final shape.

The chain code representation of a shape has a number of interesting properties. First, in order to increase the size of the shape represented, one simply needs to copy each number in the string of vectors. For example, using the scheme shown in Figure 2.1a, a rectangle can be represented by the string: 0246. In order to double the side length of the square, one changes this string to: 00224466. Another interesting property is that to rotate a shape by the smallest angle allowed, one just adds the number 1 to the chain code. So a rectangle 022466 would rotate to: 133577. This rotation actually illustrates one of the flaws of this method of shape representation: without a predefined condition, it is not rotationally invariant. In addition, one must determine where on the contour of the shape the starting point will be. The issues can be addressed by utilizing some specific conventions [29, 30].

Despite these problems, chain code remains a popular method of shape representation, and has been improved upon in a number of ways since its first implementation, including basing the chain code on enclosed cell vertices to draw relations between enclosed area and the chain code [31], creating histograms of chain codes to generate invariant features [32], combined with invariant-moment descriptors for improved image registration [33], among others. Also, a wide variety of shape information can be extracted from chain codes, including other shape descriptors, such as moments [34–37].

Bribiesca and Guzman [38] suggest the use of the derivative chain code, where two successive links are represented by a 2, a convex corner by a 1, and a concave corner by a three. If an arbitrary starting point is chosen and a shape code generated for a given object, one can then find a unique representation, invariant to rotation, by identifying all cyclic permutations of the chain code and selecting the one which has the lowest value when read as an integer. There are also a number of different methods to create similarity measures for chain codes, three of which are detailed in Mehtre *et al.* [39]

Overall, chain codes offer a method of shape representation which satisfies several of the criteria laid out in Section 2.1. They are applicable to all two-dimensional shapes, are capable of being rendered translation, rotation, and scale invariant, and have a similarity

measure. However, these descriptors fail to map visual shape information in an intuitive way, and also lack a hierarchical structure as discussed in Section 2.1.

A different approach to contour structural shape description is that of Grosky and Mehrotra [3] and Grosky *et al.* [40]. Inspired by database management systems, the primitives used are the vertices of polygons which are matched to shape boundaries using polygon approximation [41]. For each vertex, the internal angle, distance to the next clockwise vertex point, and the x and y coordinates are used to create a four-dimensional feature vector. For a given set of object, five features are selected, which are usually those with the sharpest angles, as these are considered most representative of the given shape.

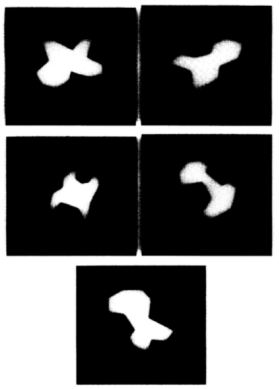

(a) Input data for feature index tree.

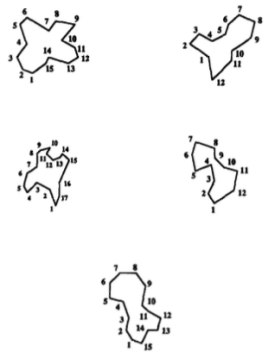

(b) Polygonal approximations of shapes input data.

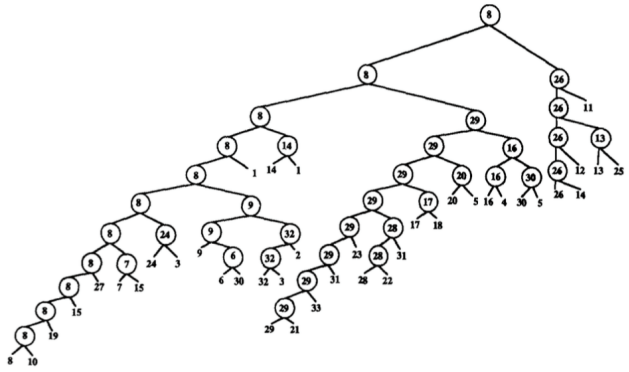

(c) Feature tree generated by features gathered from polygonal approximations.

Figure 2.2: Input data, features, and feature index tree from [3]

These sets of features are then converted into *feature index trees*, which are binary or m-nary trees with features as the leaf nodes. Each interior node contains a feature which is either an observed feature or an artificial feature and the members of a subtree are more similar to the root of that subtree than to the root of the sibling subtree. The entire process

is shown in Figure 2.2.

Object identification is performed in two stages. The first involves feature matching. In this stage, a test feature is selected from the object to be identified and is compared to those in the feature index tree. Starting with the first entry in the tree, the algorithm follows branches of the tree, checking if the similarity of the test feature to that of the current node. If the similarity is above a certain threshold, all shapes associated with that feature are returned. In the next stage, the full set of object features is compared to those of the returned shape, and the shape which minimizes some distance function is returned as the matched shape.

The distance function $d(s_1, s_2)$ for two strings of features (s_i) is measured as the cost of converting one feature vector to another. This is detailed in [42] and is based on four separate cost calculations: removing the corners of s_1, inserting corners of s_2 into s_1, changing corners of s_1 to corresponding corners of s_2, and smoothing corners in either string. Except for the smoothing, the other costs are based on Levenshtein distance [43]. Despite the fact that the Levenshtein distance is a metric, the cost function $d(s_1, s_2)$ is not because it is not symmetric.

Because it was inspired by and developed as a database searching technique for shapes, the polygonal approximation and feature index tree approach is only capable of identifying shapes which have previously been identified and testing of this technique was only done through object identification. However, that does not mean that it cannot identify the most similar shapes through its similarity measure. A potential source of error in this method is the polygonal approximation. The quality of a polygonal representation of a shape is completely dependent on the number of vertices defined by the approximation algorithm. While a number of algorithms exist which can minimize errors due to the approximation, there can never be a completely accurate polygonal approximation, because this would consist of all points on the boundary. Therefore there will always be a trade off between having a string of vertex vectors which is of a workable length and accurately representing the actual, observed boundary.

Taken altogether, the system devised by Grosky and Mehrotra, combines a shape descriptor: strings of feature vectors; with a system to identify and match similar shapes:

the feature index tree and associated search methods. The descriptors themselves satisfy a number of the previously discussed shape descriptor criteria of Section 2.1. This method can be applied to any two dimensional shape which can be described through polygonal approximation, although there are limits. Shapes with holes cannot be described using this method, because it can only be calculated for object boundaries. These descriptors are starting point invariant, but because the feature vectors include x,y coordinates and distance to the next point, the total descriptor string is not invariant to translation, rotation or scaling. The string structure of these descriptors also means that they cannot not have a hierarchical structure, so it is only useful for determining shape similarity. Extracting additional information is difficult because of this as well.

Both the chain code and Grosky's method have the primary benefit of structural methods, which is that it is much easier to identify occluded shapes. This is due to the string nature of the shape descriptors. If a shape is described as a string of identifiers which represent boundary points, an occluded shape will not have an identical string, but a consecutive subset of points will still be present, unless the entire shape is completely obscured.

These methods do have a number of drawbacks. Because they only capture boundary information, they cannot capture information regarding the bulk of the shape, especially any holes. In addition, while information can be extracted from the strings of features, it is usually difficult to get intuitive information regarding the nature of the shape through simple observation of the descriptor. Closely related to these issues is the fact that these descriptors cannot capture any global information. It may be possible for a subset, or a large number of subsets of a feature string to match between two shapes, which would indicate similar or identical shapes, however, it is still possible for the overall nature of the two objects to remain dissimilar. This is because the string of features approach is only capable of looking at features individually: it can match them, or groups of them, but it cannot get an overall view of the nature of the original object without using an another method to extract this information from the boundary. Also, the collection of features can be sensitive to noise.

Finally, the matching presents its own set of difficulties. In order to make a successful match, the measured features must already be stored in a database. Therefore, in order to

know what features should be focused on in the observed object, it is necessary to know what features are already stored. It also presents difficulty in building a shape database, because it is impossible to gather every possible useful feature for every class of shape. The difficulty in matching is compounded by the fact that many times the distance measure is not necessarily a metric, which makes interpretation of similarity difficult.

2.2.2 Contour-based global descriptors

Similar to contour-based structural descriptors, contour-based global descriptors utilize the perimeter of the object in question to quantify the shape. However, unlike the structural descriptors, these values capture information regarding the entire boundary in a single value. A number of relatively simple forms of these descriptors are common knowledge, including perimeter, compactness, and convexity.

Perimeter is a well known measure of the total length of the object boundary. Compactness, C, is typically measured as P^2/A. Because the circle is the two dimensional shape which contains the largest area, when compactness is measured, a comparison is being made to a circle with equal area. The compactness for a circle is equal to 4π, so commonly compactness is normalized by this value and called "circularity". Much more information on compactness and circularity can be read in Montero and Bribiesca [44]. Convexity is frequently given as: $c = \frac{P_{ch}}{P_s}$, where P_{ch} is the perimeter of the convex hull of the shape in question and p_s is the perimeter of the shape in question. Convexity can be used to measure the relative smoothness of the perimeter of the given object. A value of 1 indicates that the shape is perfectly convex, while a lower value indicates increasing concavity.

A commonly referenced measure of shape similarity which is calculated using the object boundary is the Hausdorff distance [45]. However, this measure is not a shape descriptor, and can only be used to compare sets of points. For two finite sets of points $A = \{a_1...a_p\}$ and $B = \{b_1...b_q\}$, the Hausdorff distance is defined as:

$$H(A, B) = \max(h(A, B), h(B, A)) \tag{2.2}$$

where

$$h(A, B) = \max_{a \in A} \min_{b \in B} ||a - b|| \tag{2.3}$$

and $|| \cdot ||$ represents the norm on the points of A and B (in this case the Euclidean norm). $h(A, B)$ is known as the directed Hausdorff distance. For each point in A it finds the closest point in B, and the distance between them, and then returns the largest distance. Effectively, it gives the distance between the point a which is farthest from the nearest point b. The Hausdorff distance calculates this, both $h(A, B)$ and $h(B, A)$ and returns the value which is the largest. It therefore finds the maximum distance between nearest neighbors of the other set.

Hausdorff distance can therefore be taken as a measure of the similarity of a set of points. If these points are the boundaries of two objects, one can use this as a measure of the similarities of the objects. Hausdorff distance has been used to for object matching in a variety of applications [46–48], but in all cases, some adjustment has to be made to make up for the sensitivity of this method to basic transformations, such as translation, rotation, and scaling, as well as to noise.

A more complex set of descriptors of this type is are what are known as *shape signatures*. Sometimes referred to as *time series*, these are one dimensional functions or signals that represent a particular measure as the boundary is traveled. Measured values include centroidal profile, centroid distance (Figure 2.3), and cumulative angle, just to name several [49]. These functions, regardless of measured quantity, are necessarily periodic for a closed object, so rotational invariance can be achieved by translating the function to some arbitrary but pre-defined orientation, which is equivalent to moving the starting point to a pre-defined location. Scale invariance can be achieved through a variety of methods depending on the measure being tracked. For angular measurements, no adjustment is necessary, but for distance-based functions, it is required. In order to calculate the centroid distance shape signature shown in Figure 2.3, the centroid is first identified, and the boundary found through edge detection. For each point along the boundary, the distance to the centroid is found, and plotted. This type of one dimensional function is a simple concept but can be used in a number of interesting ways to identify relevant shape information. Usually,

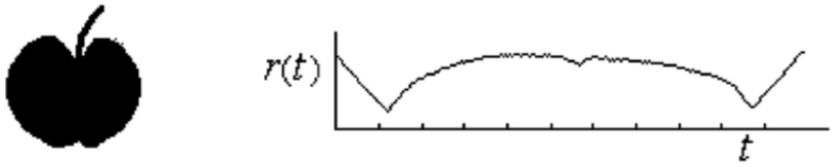

Figure 2.3: Apple silhouette and resulting centroid distance shape signature. From [4]

shape similarity is then computed by some distance measurement that is computed between two shape signatures. The majority of these in the literature are the Euclidean distance, dynamic time warping [50], and Longest Common Subsequence [51].

Traditionally, rotational invariance is achieved in two major ways: brute force matching, and landmarking. Landmarking is based upon finding one (preferably) or more features from which to start the time series conversion. Landmarking can be effective, but the main problem is that it requires feature extraction prior to the actual analysis. This usually requires a deep enough prior knowledge of the class of shapes being examined to identify these features ahead of time. In addition, it could be the case that damage to a real-world shape could change the features present at the boundary, and eliminate these previously determined locations. In some cases, the shape may be oriented to its primary axis, but this is only effective in shapes which have well defined axes. Additionally, primary axis detection can be sensitive to noise or, once again, to damage.

With brute force matching, the problem of rotational invariance is solved by circularly rotating the shape signature to all possible orientations, and subsequently measuring the similarity to each other signature. While this has shown to be the most effective method of achieving rotational invariance [52, 53], it is very computationally expensive. Its success may be due to the fact that while other methods try to extract information and match accordingly, which could result in the loss of information, the brute force method uses the entire shape profile and therefore no information is lost [5].

In Keogh *et al.* [5], the researchers present a method for implementing fast rotational invariance and demonstrate its effectiveness on a number of real world problems. The method is primarily based around *early abandonment* and a hierarchical shape signature comparison approach. When utilizing early abandonment and using some threshold, r,

during the computation of the Euclidean distance, if the current sum of distances between the two sets of data points exceeds r^2 then the actual distance between the two signatures will be less than r and so matching can cease.

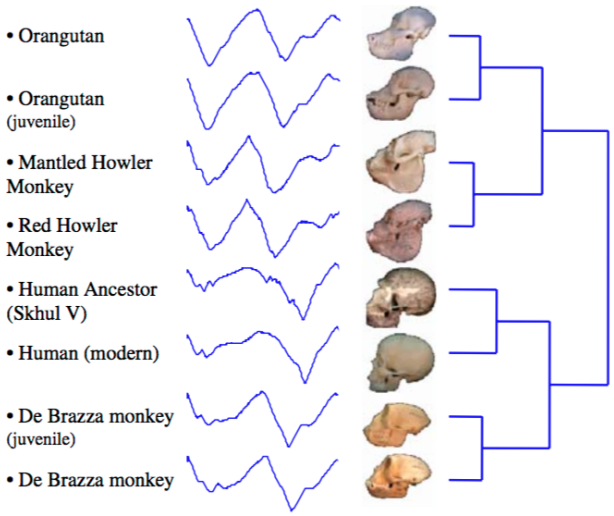

Figure 2.4: Ape skulls, resulting shape signatures, and resulting hierarchy of similarities. From [5]

In addition to the early abandon, Keogh *et al.* implemented a method of shape signature storage which allows rapid comparisons between shapes to be matched and those already stored. This is combined with a function known as LB_Keogh. While the details of this function are beyond the scope of this work, it allows for the quick comparison of stored shape signatures to those which are being tested. Utilizing the entire method, it is possible to quickly draw conclusions regarding the similarity of shapes, and even construct hierarchies of similarity as shown in Figure 2.4. From observation it is clear that the quality of the similarity matching is good, which was achieved using only the Euclidean distance between time series.

The shape signature techniques described here generally satisfy a large number of the shape criteria laid out in Section 2.1. They are applicable to any 2D shape, can be rendered invariant to translations, rotations and scalings, and also provide a similarity measure with a wide degree of discrimination ability. While not all methods are insensitive to noise, implementation of dynamic time warping and longest common subsequence have largely eliminated these issues.

Another commonly used method for global boundary quantification are *Fourier descriptors* (FDs). As described by Zahn and Roskies [54], these descriptors are calculated by first considering the closed and clockwise boundary of some shape, $Z(l)$ as a parametric curve, where l is the length along the boundary and $0 \leq l \leq L$. If the angular direction of the curve at a location is given by the function $\theta(l)$, then the cumulative angular function, the total

angular change between l and 0 is given by: $\phi(l) = \theta(l) - \theta(0)$. Finally, $\phi^*(t) = \phi(\frac{LT}{2\pi}) + t$, where $0 \leq t \leq 2\pi$. $\theta^*(t)$ is invariant to rotation, translation, and scaling of the original boundary. If it is expanded into a Fourier series:

$$\phi^*(t) = \mu_0 + \sum_{k=1}^{\infty} A_k \cos(kt - \alpha_k) \tag{2.4}$$

then every set of $\{A_k, \alpha_k\}$ represents a single curve and there is no redundancy in the information contained. For these FDs, some disadvantages are that some sequences of $\{A_k, \alpha_k\}$ do not represent closed curves, that for discontinuities generated by shape boundaries with sharp angles, A_k will decrease slowly, and that reconstruction requires numerical integration.

The FDs described by Granlund [55] are defined differently, which slightly changes their properties. First, again the clockwise shape boundary is defined as a parametric curve $Z(l)$. A point moving around the boundary generates a complex function $u(l) = x(l) + iy(l)$ which is periodic with a period of L. The FDs are now:

$$a_n = \frac{1}{L} \int_0^L u(l) e^{in(2\pi/L)nl} dl \tag{2.5}$$

with

$$u(l) = \sum_{-\infty}^{\infty} a_n e^{jn(2\pi/L)l} \tag{2.6}$$

With these FDs, the advantages are that for a sequence of a_n for which (2.6) converges and describes a closed curve, $|a_n|$ decreases quickly as $n \to \infty$, and reconstruction is relatively simple. However because of the definition of $u(l)$, there are some restrictions on a_n.

Using both the FDs described here, Persoon and Fu [56] matched handwritten symbols from (ϕ,1-9), which were stored as 24x24 binary images. There are a total of 49 x 3 x 10 of these numerals, each handwritten by a different person. Two sets of these were used for training and another for testing. Distance between two FDs is calculated as:

$$d(\alpha, \beta) = [\sum_{n=-M}^{M} |a_n - b_n|^2]^{1/2} \tag{2.7}$$

with $n \neq 0$ in order to make the values translation invariant. In addition to handwritten numerals, a series of machine parts were photographed, and the images binarized into images of 60 x 60 pixels.

In the case of the handwritten characters, some preprocessing was necessary. Each numeral was normalized for scaling rotation, position, and descriptor starting point. Utilizing all the harmonics up to a value of $n = 15$ led to an error rate 10.6%. In the case of the machine parts, the training set consisted of 10 different parts, with only one orientation, scale, and position recorded. Machine parts were correctly assigned to proper classes.

Fourier descriptors have been utilized in a variety of different studies and have been modified and improved upon in a number of ways [57–60]. There are several advantages to these descriptors. The first is that they allow reconstruction of the boundary in question, which means that they contain all of the shape information of the original outline. In addition, this feature means the nature of the shape contained in each set of the series of descriptors can be determined by performing a reconstruction using only up to that value of n. These descriptors are also easy to compute, simple to normalize, and the many different values of n mean that both local and global features are captured. When measured against the criteria in Section 2.1, properly implemented FDs fulfill nearly all of the criteria and provide a good method of shape quantification overall.

Global boundary descriptors, including shape signatures and Fourier descriptors, provide a number of advantages in the evaluation and quantification of shape. Because they are global descriptors, they can quantify the overall nature of a boundary. Additionally, with these descriptors, the fineness or coarseness of an analysis can be modified to examine more global or local information at will. They also are more intuitive to work with and in many cases relate well to visual shape information that can be observed by the human eye. Because the features are global, noise cannot influence the quantified features to the degree that it can with local features.

Despite the fact that these descriptors are generally quite effective in measuring shape, they have a few downsides. First among these is that they do not apply well to occluded shapes. While partial matching can ameliorate this to some degree, it typically involves the use of more and more local features which must be identified by some other means, which

has its own potential for error.

2.2.3 Region-based structural descriptors

Region-based structural descriptors function similarly to boundary-based structural descriptors. The primary difference between the two is region-based descriptors utilize shape information from the entire bounded area, whereas boundary-based descriptors limit information to the edge of the shape. Region-based descriptors can benefit from an increase in total available shape information and often times more accurately describe shape that contain voids. However, methods of breaking up a region into its components for use in a structural descriptor are more unclear than methods to break up a one dimensional boundary.

The *medial axis transform* (MAT) is a commonly utilized descriptor of this type, and also provides an example of what Loncaric calls *space-domain* descriptors, which, instead of returning a scalar value, returns another image. The previously mentioned *shape signature* descriptors also fit this description [28]. The MAT was first introduced by Blum as a method to describe shape which could help further understand how shape is perceived by organisms [61].

Blum illustrates the formation of a medial axis construction by first considering only the boundary of a given shape. If this shape is dropped into a field which propagates waves, not unlike a pool of water, the waves move inward. Where these waves meet, lines are drawn, which represent the medial axis of the shape. This process is shown in Figure 2.5. A transform to generate these axes proposed by Blum and now referred to as the *grassfire transform*, does so by assigning each pixel in the contained region the minimum distance to the boundary. The most intense values within the region make up the medial axis. The actual implication of this process is that each point on the line of the medial axis is the center of a circle which is tangent to the given surface at two or more locations.

As long as holes are included in the set of boundaries being examined, the grassfire transform, and thus the medial axes, are capable of reflecting these inconsistencies in the bounded region. A *radius function* defines a radius of a circle for every point along the skeleton, although Calabi refers to this as a *quench function* [62]. When combined, the

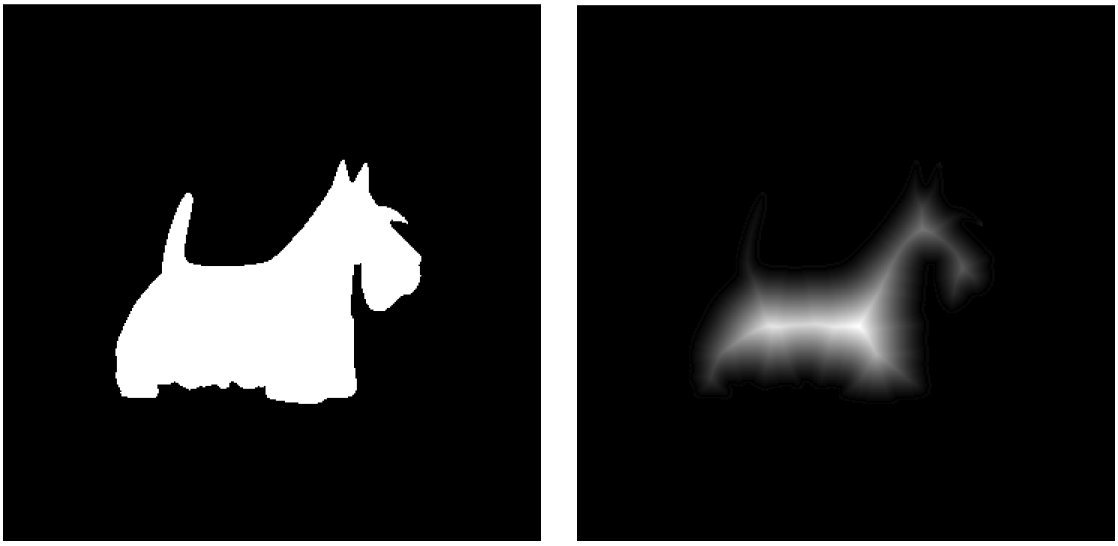

(a) Original binary image of silhouette of Scottish terrier.

(b) Grassfire transform of Figure 2.5a. The brightest lines represent the medial axes of the shape.

Figure 2.5: Formation of medial axes using grassfire transform

radius function and medial axes are capable of complete recreation of the shape in question [63]. The medial axis representation is related to the idea of *shock graphs* which have been used to compare shapes in a number of ways [6, 64–67]. The basis of this technique is shocks, which effectively label regions of the medial axis transform of an object by the way in which the radius function behaves over that section of the axis. More specifically, where the medial axes were created in a thought experiment in which the medial axes represented the meeting of wavefronts, shocks are literally treated as these meetings of wavefronts: they have a position((x, y), a time of formation, related to the size of the radius function (distance to the generating portion of the boundary), and a direction. In Siddiqi, shocks are also given a label based on the properties of the radius function on their location: a 1-shock is where the radius function varies monotonically as in the case of a protrusion. A 2-shock describes an axis where the radius function reaches a local minimum as in the case of a neck. A 3-shock describes a radius function of constant value. A 4-shock describes an axis where the radius function reaches a local maximum. Utilizing shocks allows the problem of shape similarity to be accomplished through graph matching. First, each set of connected shocks for both the model and unmatched objects are deconstructed into a shock tree. This allows the for a hierarchical view of the set of shocks which make up the object. The trees

are generated in a such a way that the shock tree has as few tiers as possible. Next, the distance between matching branches and nodes is computed by first following the two branches which are closest to each other. Distance between nodes is a function of the four properties discussed in the previous paragraph: position (x, y), time of formation, and direction. Branches are matched until they reach a termination. The use of the trees ensures a matching that is based on the topological features of each graph while also requiring a distance to be calculated as well.

Overall, matching using these features provides some distinct advantages. Medial axes and the matching provided by shock graphs allow topological information about the constituent parts of the shape to be preserved. This allows for *classes* of shapes to be matched, even if they have been distorted to a degree. In Figure 2.6 from Sebastian *et al.* [6] for example, missing regions of the jet are picked up by the descriptors, but due to the topological and structural nature of the features, a match remains possible. This is true for the occluded shapes as well, primarily because features remain which can be matched.

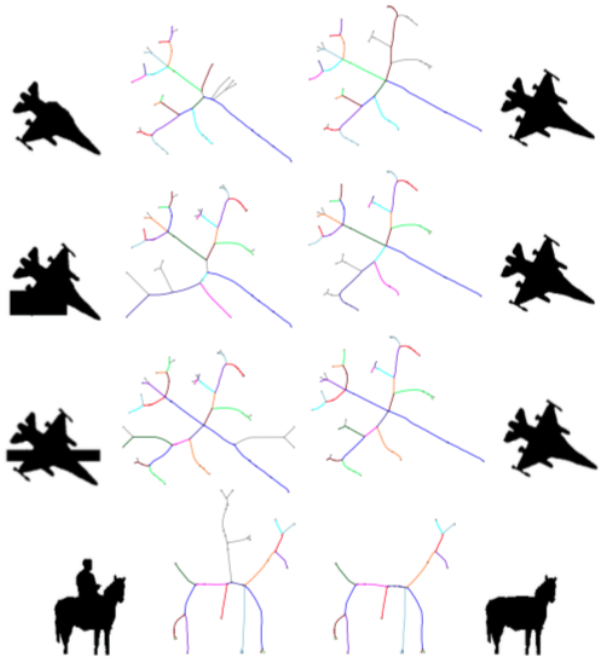

Figure 2.6: Matching shock graphs despite occlusion or missing regions. Same color shocks represent matches. From [6].

These descriptors, which use shape boundaries to gather structural information regarding the bulk object, are sensitive to boundary noise. Distortions along the boundary can result in medial axes which do not match the reality of the shape. While this should not affect the matching, as can be seen from above, it still alters the measured shape.

Region-based structural descriptors, as represented by medial axis transforms and shock graphs, provide a relatively unique way to examine shape. By breaking up the shape into

components, one extracts a more topological view of shape than one can get from other methods. In addition, because these descriptors are structural, they are less affected by occlusion or missing regions than others. They are hierarchical in that one can get coarse information from the axes such as the primary axis or the aspect ratio, or one can go to more detail and completely reconstruct the shape. Another advantage that these descriptors offer is that they very clearly reflect visual aspects of the shape because they return another image.

However, the fact that the information is returned in the form of an image often leads to more complicated comparison methods such as the previously discussed shock trees. Sensitivity to noise is another issue that requires relatively sophisticated methods to resolve.

2.2.4 Region-based global descriptors

Region-based global descriptors capture information on the morphology of the object by using the entire area contained within the object boundary. Similar to boundary-based global techniques, these are more suited toward object reconstruction, and tend to be more robust to noise. Moment based descriptors such as Cartesian and Zernike moment invariants are of this type, but will be discussed in more detail in Chapter III.

The grid-based method of Lu and Sajjanhar [68] is a non-moment invariant based way to obtain global shape information from a 2D object. This technique is based upon representing a shape by a series of 1s and 0s. Each is a point on a grid which is either filled and represented by a 1, or empty and represented by a 0. The resulting series of digits is treated as a vector, and similarity measured through comparing distances between these.

The size of the grid is selected so that it just covers the shape in question. The grid cells are square and in the case of the work by Lu, set to be 24x24 pixels. However the authors recommend between 10 and 20 pixels to a side as a good compromise between accuracy and computation time. Next, any cell which is covered at least 15% by the object is considered filled and represented by a 1. The cells are moved along, left to right, top to bottom, and turned into a string of 1s and 0s.

While the grid method is inherently translation invariant, the shapes have to be normalized beforehand to render them rotation and scale invariant. The normalization methods

are based around finding the major axis of the shape in question. In this case, the major axis is defined as the line connecting the two most distant points on the boundary of the object. Scale normalization is achieved by scaling the major axis to a fixed length. In the case of Lu's work, this was to 192 pixels. Rotational invariance is achieved by aligning the major axis so that it is parallel with the x-axis of the grid. There are actually two possible orientations which are 180° apart from each other. In order to address this, both orientations are stored to be matched against.

The size of the grid is based upon the shape being examined. Because all shapes should have the same length, the number of cells in the x direction can be fixed. However, the number of cells in the y direction is based on the size of the minor axis. The minor axis is perpendicular to the major axis and just long enough so that a rectangle constructed from the major and minor axes would completely cover the shape. Therefore, the lengths of the binary strings will change based upon the aspect ratios of the shapes being examined.

The distance between these strings, and therefore the similarity between shapes, is then based upon simple calculations. If the string lengths are very different, the shapes are assumed to have a different aspect ratio and therefore be different shapes. If however, the string lengths are within some predefined threshold, in this case within 24 characters, the shorter string is padded with zeros and then the similarity calculated as if they were the same length. This distance calculation is simply the total number of unlike positions between the strings.

This method proved effective in shape recall for a number of different geometries and outperformed a Fourier descriptor (FD) based method (previously discussed). Lu suggests this is because these descriptors are much less sensitive to noise, as any small changes to pixels around the object boundary will push an individual cell over or under the 15% threshold.

This grid-based approach has been examined in several different works [69–71]. Chakrabarti *et al.* improved the method with an adaptive resolution (AR) approach. AR increases the resolution of the grid used to capture shape where it improves the quality of the representation, but reduces it where it does not. This allows for a finer detail where necessary while also increasing the efficiency of the representation.

Grid-based descriptors are generally insensitive to noise and provide a very simple similarity measure and, as such, satisfy several of the criteria set out in Section 2.1. They can be applied to any two dimensional shape, and use a similarity measure which allows for the determination of degrees of similarity. However, they also lack several of the benefits that have been observed as properties of other descriptors. First, unless using the adaptive resolution approach, resolution is fixed, so the level of detail is dependent on grid size. Rotation invariance is based on implementing multiple comparisons to transformed model shapes, which may necessitate the storage of additional reference information. Also, the strings that are generated using grid-based descriptors do not necessarily correlate to any kind of perceptual information unless used to reconstruct the data in some way. Finally, the resulting data isn't hierarchical, meaning that resulting detail is limited to that of the data initially collected.

These drawbacks are not present in other region-based global descriptors, like moment invariants which are discussed in Chapter III. Overall, descriptors of this type tend to be more hierarchical, and, at least at the more general level of description, tend to reflect general and intuitive shape information such as circularity or aspect ratio. These shape measures are also relatively insensitive to noise. To a certain degree, they cannot replicate the kind of topological preservation of information that is present in structural descriptors, but the hierarchical nature does allow for the containment of very specific information in each aspect of the descriptor. For these reasons, region-based global descriptors among the most popular in the shape analysis community.

2.2.5 Summary

For overview purposes, descriptors of two-dimensional shapes have been classified as boundary-based and region-based, with each further classified as structural or global in focus. Boundary-based descriptors can be used to draw very specific comparisons between two objects, due to the very simple ways in which a one-dimensional boundary can be represented and compared. Region-based techniques provide a more complete picture of the object as a whole and are less sensitive to noise. Topology can also be preserved. Structural techniques allow a preservation of the local information of the object, and given that a shape is represented

using string of features, similarity measure and shape representation is simpler. Global methods preserve information regarding the entire shape as a whole, which can be lost otherwise. They are less sensitive to noise. Depending on the application, the inherent advantages and shortcomings of each approach lends it to be a more or less suitable shape descriptor of two-dimensional objects.

2.3 Three-dimensional shape descriptors

Three-dimensional (3D) shapes are more difficult to quantify than their two-dimensional (2D) counterparts. The previously established categories begin to break down in 3D. Boundary based descriptors are less useful in this case because instead of simple one-dimensional lines, the boundaries are 2D surfaces, which are much more difficult to compare to one another. Structural descriptors also begin to be more difficult to calculate because criteria for feature selection becomes more complex.

In a 2008 survey of 3D shape retrieval methods, Tangelder and Veltkamp categorized the most common 3D shape representations into three groups: feature based, graph based and geometry based [72]. In the interest of brevity, an example here will be given of each, and the advantages and disadvantages discussed.The discussion is meant to highlight a representative slice of the types of 3D shape quantifiers that are currently in use.

2.3.1 Feature based methods

Feature based three-dimensional shape representations either use extracted local features, similar to the structural methods from Subsection 2.2, or global features, which include 3D moment based techniques which will be discussed in more detail in Chapter III.

In Sipiran *et al.* [73], a method is put forth which combines global and local feature extraction in order to create an effective object retrieval system. The method is primarily built around three steps: *interest point detection*, *clustering of interest points*, *cluster-based partition*. Points on the surface of an object are automatically selected as being of interest and then clustered. The entire object is then represented as a collection of these clusters which is used to calculate a similarity value.

Automated region of interest location is found through use of the Harris 3D method,

which was developed by Sipiran and Bustos [74] and based upon the Harris operator [75]. For a vertex point on a 3D mesh, the Harris 3D operator at that point is found by first fitting a plane to the surface of the object. Using principal component analysis (PCA) the points are rotated so that their z-axis lies on the normal of this plane. A quadratic surface is then fit to the set of transformed points. This allows for derivatives of the surface to be calculated, and using these values, the Harris operator can be found at this point. Selected interest points are those that represent a local maximum for this operator.

There are three parameters which determine how the interest point clustering is performed: S: the intra-cluster constraint; R: the inter-cluster constraint; and N_m: the minimum number of points per cluster. The set of selected interest points are then reviewed and assigned to clusters in such a way that S defines the maximum distance between a point and its cluster centroid, and R defines the minimum cluster centroid distance. Finally, if after clustering, the number of points in a cluster is less than N_m, the cluster is discarded.

Partitioning occurs for each cluster by forming the smallest sphere with a center on the cluster centroid and containing all points within that cluster. The region of the mesh contained within that sphere is then extracted and used as a feature for the object. The distance between feature lists is found by comparing individual components of partitions and determining if they match. Sipiran *et al.* were able to improve retrieval effectiveness in object classes which had common distinctive parts, but where common distinctive parts were not present, little improvement was seen.

Feature based methods, including global descriptors tend to be the most commonly used in 3D object recognition because they correlate to visual shape information, are generally easy to understand, and are simple to calculate. While not obvious from the previous example, these values, especially global ones, are also more numerical in nature, making them easier to work with when compared with the graphed based methods discussed below. However, by extracting pieces of information regarding morphology and examining them outside of the context of the morphology of the shape, these methods can fail to fully capture the topological nature of the object in question.

2.3.2 Graph based methods

Graph based methods include model graph [76] and skeleton based [77] methods. These techniques are effective in preserving topological information and, as has been shown in two dimensions with medial axes and shock graph approaches, allow effective matching of classes of shapes even in the presence of distortion. However, extending these to three dimensions can be difficult due to a high computational cost and sensitivity to noise on the boundaries [7].

To address these problems, Shinagawa *et al.* utilized both Morse theory [78] and work by Reeb [79] to create what they called the *extended Morse method*, which allows for the encoding of 3D surfaces as two dimensional graphs. This method, now commonly referred to *Reeb graphs* is considered the most popular graph based approach to 3D shape representation.

In order to generate a Reeb graph, a continuous function, μ, is defined on a manifold. A simple and common definition for μ is the z position. From this, two points on the manifold, x_1 and x_2, are considered equivalent if $\mu(x_1) = \mu(x_2)$ and if x_1 and x_2 are in the same connected component of $\mu^{-1}(\mu(x_1))$. An example of this is shown in Figure 2.7, where μ is the height of the point on the torus surface, represented by color.

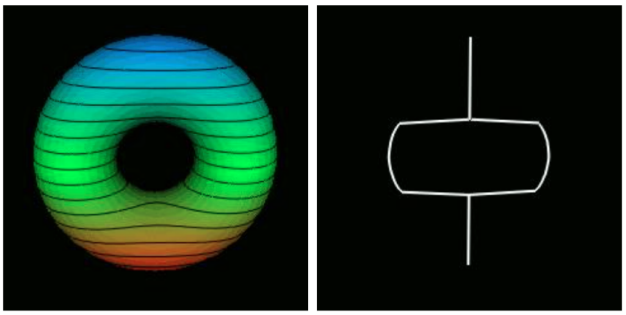

Figure 2.7: Example of Reeb graphed based on height along z axis for a torus. From [7].

Because height is not rotationally invariant, Hilaga *et al.* [7] proposed the utilization of geodesic distance as μ. However, because a constant source point to serve as a reference for this distance could prove sensitive to small deformations, $\mu(x)$ was defined as the sum of geodesic distances from x to all other points on the surface, where x represents a given vertex of a mesh. This value therefore represents the degree to which the object is near the center of a surface and is stable across shapes within the same class.

Hilaga *et al.* also introduced the idea of the *multi-resolutional Reeb graph* (MRG). This

graph is actually a collection of Reeb graphs which are calculated for increasingly coarse thresholds of $\mu(x)$. First, a Reeb graph is constructed using a fine threshold for similar μ values. Values that are within this threshold are considered equivalent and generate an *R-node* which is simply a representation of a collection of equivalent vertices, which are connected by lines representing topological connection. The combination of R-nodes and lines constitutes a Reeb graph. The threshold is increased, the R-nodes recalculated, and this is repeated until there is only one node remaining, representing the entire shape. The result is a series of Reeb graphs which represent the topological nature of the shape at a variety of resolutions.

Similarity between shapes is determined by moving through these increasingly fine representations and determining the similarity between nodes by calculating a similarity measure that is based upon the normalized area of the represented mesh, and the maximum value of $\mu(x_1) - \mu(x_2)$ for each R-node. By iteratively moving from coarse to fine representations and moving node to node, the topological information of the object is preserved.

The MRG method and its variations have proven to be effective in matching deformed and transformed shapes as well as shapes within classes. Hilaga *et al.* showed that the method is effective in matching similar shapes in a 230 object database that only included different shapes, not those that were transformed.

Bespalov *et al.* used three separate test sets of 3D objects to examine the effectiveness of the MRG method. The method proved effective in determining similar shapes on a set of cubes with holes, a set of brick-shaped models with holes, a set of basic primitives (spheres, cylinders, cubes, and tori), and on complicated CAD models. The method performs less well on 3D modeled LEGO parts, where plate shapes and axles with x-shaped cross sections were conflated [80]. The Reeb graph method has also been extended to clouds of points for body scanning [81] and to partial object matching for rigid and non-rigid transformations [82], among others.

Graph based methods provide a level of topological information that cannot be matched by other methods, but lack the global information that can be determined from feature based methods. Overall, these methods can be applied to any 3D shape, can be made invariant to rigid 3D transformations as well as robust to some non-rigid ones, and provide a hierarchical

view of the shape, especially with the MRG based techniques. However, in many cases, the way in which the measured similarity relates to perceptual shape information is not obvious.

2.3.3 Geometry based methods

Geometry based methods include view based techniques, volumetric error, and weighted point set. These methods generally analyze shape similarity as a series of transformations. The fewer the number of transformations required to make one shape identical to another, the more similar the two objects are. The volumetric error method described by Sánches-Cruz and Bribiesca [83] views each object as a collection of voxels and transformations are the required movements of these. Besl and McKay's weighted point set method [84] represents a shape as a collection of points, polylines, parametric curves, or some other representation, and determines the transformations which will result in the minimum distance between a measured shape and a model shape.

View based techniques attempt to emulate how shape is perceived by representing a 3D shape as a number of 2D projects or views. If two objects are similar, they should have similar 2D projections. Funkhouser *et al.* developed a method of sketch queries which allows a user to draw an object, and similar 3D models will be returned [85]. A similar method was utilized by Nayar and Murase [86], although many more views of the 3D object were employed.

In the sketch queries method, thirteen contours are generated for each 3D object from a set of thirteen orthographic projections. These views are compared to the user sketched image using a similarity metric, and the model with the highest total similarity using every possible view is selected as most similar. As the sketches cannot be guaranteed to be closed contours, the Hausdorff distance or other similar metrics can be used. Ultimately, the sketch method was found to be complimentary to text-based image searches, with text being most discriminating across classes while the sketches improved precision within them [85].

Blanz *et al.* utilized the view based method with support vector machines (SVM) to identify 3D models of chairs and achieved an error rate as low as 1%. Roobaert and Hulle found that a SVM approach to shape recognition utilizing color and shape could operate effectively with as few as eight views of the object before losing retrieval power, and that

the method proved effective for unseen views [87]. Cootes *et al.* treat each view as a Markov Chain and find the maximum *a posteriori* probability which will determine the object to be returned.

Geometry based methods, particularly view based ones, are becoming much more popular as machine learning is becoming more powerful. They can be applied to any 3D shape, and similarity measures are intuitive. However, in general, the descriptors do not correlate to perceptual aspects of shape and are not hierarchical.

2.3.4 Summary

Like 2D descriptors, the variety of 3D shape representations allows for many different and nuanced approaches to shape. While feature based methods allow for a more intuitive matching process, they become difficult to calculate automatically in three dimensions. Topological connectivity can be preserved with graph based methods, but these are sensitive to noise, can miss global information, and can be computationally intensive to calculate. View based methods allow simple feature representation as a collection of 2D projections, but they do not allow global information extraction and are highly reliant on the type of matching method used. As with two-dimensional descriptors, descriptors of three-dimensional shapes are more or less reliable depending on the particular application, with any one potentially being superior depending on the context of the shape description.

CHAPTER III

Moments and moment invariants in shape classification

Moment invariants have been a significant method of shape classification since Ming-Kuei Hu published his work on seven position, orientation, and scale invariant combinations of Cartesian moments in 1962 [88]. Moments values represent a mapping of a shape onto a basis function which controls the nature of the shape information contained. Moment invariants are combinations of these moments which are invariant to simple transformations of the object and therefore contain only shape information. They provide a global shape descriptor which has the potential to satisfy a large number of the criteria of an effective shape descriptor laid out in Section 2.1. Moments can be applied to any two- or three-dimensional shape. They provide a similarity metric that, due to a hierarchical structure, can identify similar shapes while still identifying their differences, and they relate to visual shape information. Utilizing invariant combinations of these moments renders the result insensitive to changes in location, orientation, and size, and in some cases, other kinds of transformations. These factors make moment invariants one of the most commonly used shape descriptors in many applications across a number of fields.

Moments of one dimension have long been used to gain important information regarding a probability function. For any random variable X, the general statistical, or one-dimensional, moment function is:

$$m_p = \int_{-\infty}^{\infty} x^p f(x) dx \tag{3.1}$$

where $f(x)$ is the probability function of X and p is the order of the given moment.

The order controls the information that is returned by the statistical moment. Using

a value of $p = 1$, the moment represents the average value of the probability function. A value of $p = 2$ gives a moment value that represents the *variance*, or the spread of the distribution. $p = 3$ returns a moment value that is equal to the *skewness* of the function, which describes the symmetry of the distribution, with a negative value indicating left-skewed, a positive value indicating right skew, and a value of 0 indicating no skewness. A value of $p = 4$ returns the *kurtosis* of the function, which is a measure of the weight of the tail [89]. A higher value represents more influence of the tails in the variance, while lower values indicate less, which is typically the case in distributions which are bounded.

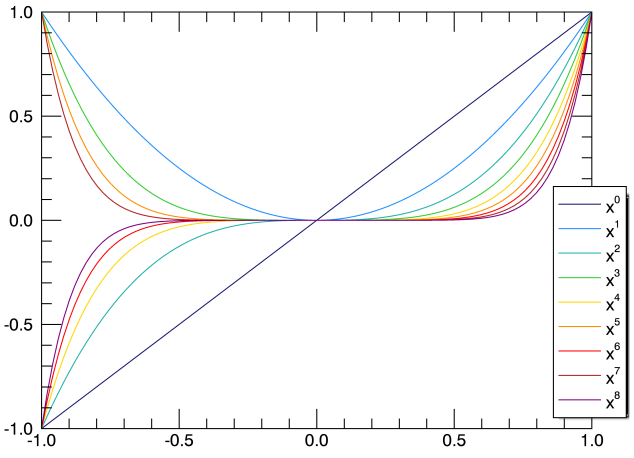

Figure 3.1: Weight function for basis function x^p with increasing values of p.

The above depends on a distribution is centered about the origin. Without this, variance, skewness, kurtosis, and statistical moments of higher orders become less meaningful because the interaction with the basis functions shown in Figure 3.1. In order to center the distribution at the origin, the general moment equation changed to:

$$m_p = \int_{-\infty}^{\infty} (x - m_1)^p P(x) dx \quad (3.2)$$

which renders any moments of order $p > 1$ invariant to the x position of a distribution. Which is to say that regardless of the location of the function, as long as it has the same shape, it will have the same set of moments as well.

With increasing order comes increasingly more complicated interpretations of the meaning of the resulting moment value. This is a problem with statistical moments, as well as for the other moments that are discussed further in this chapter. While lower order moments correspond to intuitive values, increasing the order changes the weighting of the x values in a way which makes their influence more and more subtle. This is clear from Figure 3.1, which illustrates these functions for $-1 \leq x \leq 1$. As p increases, the weight of basis func-

tions is moved progressively more and more away from the origin and the degree to which they vary moments of order $p-2$ decreases.

Due to the symmetry of functions with an even value of p, the value for any odd order of statistical moment will be positive. This same effect leads to negative or positive values for m_p. However, for a symmetric function which is centered around the origin, this also means that odd orders of m_p are equivalent to zero. This needs to be kept in mind when attempting to calculate statistical moments for functions.

The above holds true not only for continuous functions of x but for discrete functions as well. In this case, $P(x)$ can be represented as a sum of delta functions. This leads to one dimensional moment functions of the form:

$$m_p = \sum_{i=0}^{\infty} x^p P\{x = x_i\} \tag{3.3}$$

$$m_p = \sum_{i=0}^{\infty} (x - m_1)^p P\{x = x_i\} \tag{3.4}$$

If the probability of any one value of x is equal to either one or some fixed value, the discrete function begins to mirror that of a single line of pixels. Utilizing a set of moments, increasingly complex information about the distribution of these "pixels" can be gathered by collecting moments of increasing order. Making the leap from applying moments in the context of statistical analysis to the discussion of an image requires an adjustment to the function being utilized. The discrete probability function is replaced with an *indicator function*, and the limits of the summation are adjusted:

$$m_p = \sum_{i=-\infty}^{\infty} x^p f(x) \tag{3.5}$$

where

$$f(x) = \begin{cases} 1 & \text{if on relevant pixel} \\ 0 & \text{if not on relevant pixel} \end{cases} \tag{3.6}$$

By restricting the value of a pixel to either "off" or "on", it is possible to obtain relevant information from the moment values. m_0 is equal to the number of pixels which are equal

to 1 in the given system. When divided by m_0, m_1 is the average location of the "on" pixels, which is equivalent to the centroid of the system. Similar to equation 3.4, if μ_1 is subtracted from equation 3.5, one renders the distribution of pixels invariant to changes in location. With increasing order, more and more information regarding the pixel distribution is obtained, and through the use of enough moments, the distribution of pixels can be completely reconstructed.

This is exactly equivalent to the construction and use of two-dimensional moments. The transition is simple to make and only requires the addition of the other dimension and variable:

$$m_{pq} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x^p y^q f(x,y) dx dy \tag{3.7}$$

where in this case, the order of the moment is now quantified as: $p+q$. This is true for a continuous indicator function, but with digital images, it is easier to use:

$$m_{pq} = \sum_{x=0}^{X-1} \sum_{y=0}^{Y-1} x^p y^q f(x,y) \tag{3.8}$$

where X represents the number of pixels in the x direction and Y the number of pixels in the y direction.

The general structure of any 2D moment is the same: a basis function and the indicator function. For any fixed shape, the indicator function remains the same, and carries the shape information which will be encoded in the moments. The basis function determines *which* information is stored in each moment. By using different basis functions, the nature of the moments, including the quality and interpretability of the information can be changed.

3.1 Cartesian moments and invariants

Cartesian moments, sometimes referred to as *geometric moments*, utilize the basic $x^p y^q$ monomial, or $x^p y^q z^r$ in three dimensions. These moments have been the most commonly used, primarily because of their simplicity in construction and the significant amount of study they have engendered. However, the simple nature of the basis function often leads to undesirable consequences which will be discussed.

3.1.1 Two-dimensional Cartesian moments

The two-dimensional Cartesian moment equations (3.7 and 3.8) are relatively simple and encode intuitive information into the first several orders of these moments. Given these desirable attributes, they have been used for a number of different applications [90–94]. The values of the basis functions up to the second order are shown in Figure 3.2. These illustrate the relative weighting of the locations within the unit cell, and also mirror the basis functions shown in Figure 3.1, which can be thought of as a cross section through the contours of Figure 3.2 along either the x- or the y-axis.

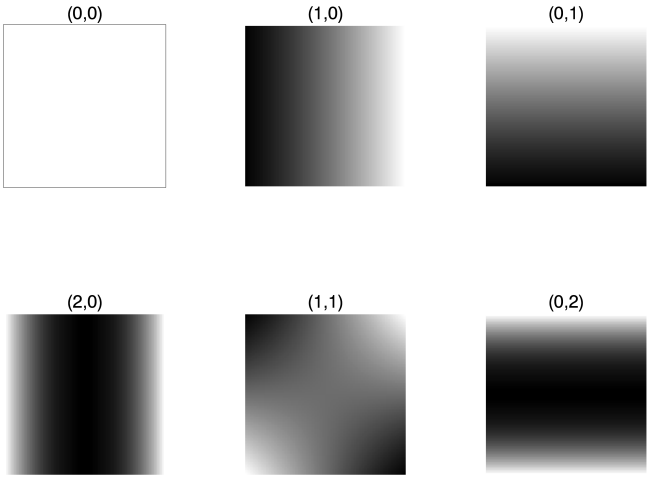

Figure 3.2: The relative values of the basis function $x^p y^q$ for the given (p, q). Values shown for $-1 \leq x \leq 1$ and $-1 \leq y \leq 1$.

It should be evident that basis functions of increasing order are not orthogonal [95]. The definition of orthogonal functions for two equations $f(x)$ and $g(x)$ with a weighting function of $w(x)$ over the interval $a \leq x \leq b$ are as follows:

$$\langle f(x)|g(x)\rangle \equiv \int_a^b f(x)g(x)w(x)dx = 0 \tag{3.9}$$

According to Teague [96], the nonorthogonality of these functions means that moments of different orders contain redundant information regarding the nature of the shape. While this does not mean that these values are useless, it means that each additional moment

added to a set adds less information than would be the case were they orthogonal.

This has implications for the reconstruction of shapes. As discussed in Chapter II, and according to Brady [24], shape descriptors are either information preserving or non-preserving. Moments are information preserving, because, with a complete set of moments, the entire object can be reconstructed [96]. However, by using a set of nonorthogonal basis functions, the reconstructions are less efficient and require more moments. As seen in Figure 3.3, sufficient detail to determine which letter is being reconstructed only occurs at approximately the 40^{th} or 50^{th} order. This can be reduced by using orthogonal basis functions. While this work does not focus on reconstruction, it provides a powerful tool for the visualization of precisely what information is stored in each moment. More information on Cartesian moment reconstruction can be found in Gorbel *et al.*

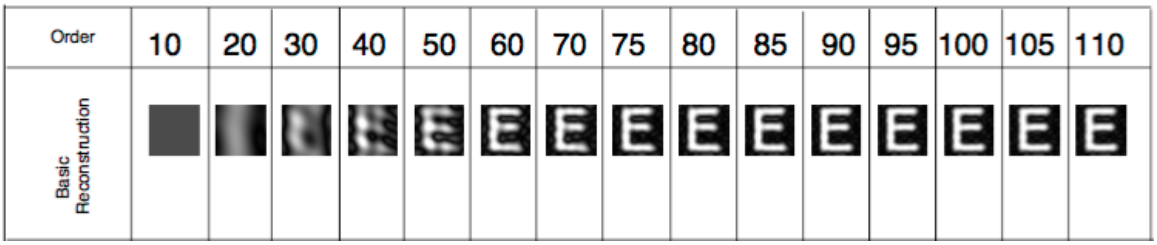

Figure 3.3: Reconstruction of the letter "E" using geometric moments. From [8].

2D moments contain information regarding the relative distribution of pixels in the image. Inserting $p = 0$ and $q = 0$ into equation 3.7 returns the definition of area. m_{10} and m_{01}, when scaled by $\frac{1}{m_{00}}$, give the x and y coordinates of the centroid of the shape in question respectively. Utilizing these values, the moments can be made invariant to two types of basic transformations which will be encountered: translation and scaling.

3.1.2 Cartesian Moment Invariants

A primary attribute of quality shape descriptors discussed in Section 2.1 is invariance to basic transformations, which include translation, rotation, and scaling. Primarily, this seeks to isolate shape from location and orientation information, which would otherwise affect the values of the moments. Simple translation invariance can be achieved through the method for 1D moments in equation 3.4. This simply involves subtracting the centroid coordinates

from equation 3.7 as follows:

$$\mu_{pq} = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}(x-\bar{x})^p(y-\bar{y})^q f(x,y)dxdy \tag{3.10}$$

where $\bar{x} = \frac{m_{10}}{m_{00}}$ and $\bar{y} = \frac{m_{01}}{m_{00}}$. μ_{pq} are known as *central moments* and are completely translation invariant. The result of the subtraction of the central moments in equation 3.10 is equivalent to shifting the centroid of the shape to the origin. By utilizing these moments in all future calculations, all resulting values are also translation independent.

A relatively simple way to attain scale independence is to utilize the 0^{th} order moment as described by Hu [88]:

$$\frac{\mu_{pq}}{\mu_{00}^{\frac{p+q}{2}+1}} \tag{3.11}$$

This normalization forces the given moment to return the same value for a shape regardless of its area. This is only true for isotropic scaling.

In addition to location and size invariance, rotational invariance was addressed by Hu as well. The first method is referred to as the method of principal axes. In this method, a unique axis of the object is determined with second order moments. The shape is then rotated so that this axis coincides with the x-axis. In this case, the principal axis is defined as the eigenvector corresponding to the largest eigenvalue of the second order central moment matrix, more commonly referred to as the *structure matrix* [97]:

$$\begin{bmatrix} \mu'_{20} & \mu'_{11} \\ \mu'_{11} & \mu'_{02} \end{bmatrix} \tag{3.12}$$

where $\mu'_{pq} = \frac{\mu_{pq}}{\mu_{00}}$.

The two resulting eigenvectors, largest and smallest, of this matrix represent the major and minor axes of the image, respectively. These are essentially the long and short axes of the object. The angle corresponding to a major axis parallel to the x-axis also corresponds to a μ_{11} value equal to zero. This angle θ is given by:

$$\theta = \frac{1}{2}\arctan\left(\frac{2\mu_{11}}{\mu_{20}-\mu_{02}}\right) \tag{3.13}$$

This angle is not necessarily unique, however, and so would be unsuitable for complete rotational invariance. Apply the additional constraints that $\mu_{20} > \mu_{02}$ and $\mu_{30} > 0$ after rotation, the angle can be uniquely determined [88].

Manually rotating the image in question can be a less than ideal solution. When performed on digital images, it can result in a loss of information or the addition of noise along the borders. Using θ, Teague suggested a virtual rotation according to the following [96]:

$$\mu_{pq}^{r} = \sum_{r=0}^{j}\sum_{s=0}^{k}(-1)^{k-s}\binom{j}{r}\binom{k}{s}(\cos\theta)^{j-r+s}(\sin\theta)^{k+r-s}(\mu_{j+k-r-s,r+s}) \tag{3.14}$$

which avoids the problems associated with the rotation of the digital image.

The second method suggested by Hu for rotation independence is a combination of central and scale invariant moments [88]:

$$\begin{aligned}
I_1 &= \mu_{20} + \mu_{02} \\
I_2 &= (\mu_{20} - \mu_{02})^2 + 4\mu_{11}^2 \\
I_3 &= (\mu_{30} - 3\mu 12)^2 + (3\mu_{21} - \mu_{03})^2 \\
I_4 &= (\mu_{30} + \mu_{12})^2 + (\mu_{21} + \mu_{03})^2 \\
I_5 &= (\mu_{30} - 3\mu_{12})(\mu_{30} + \mu_{12})[(\mu_{30} + \mu_{12})^2 - 3(\mu_{21} + \mu_{03})^2] \\
&\quad + (3\mu_{21} - \mu_{03})(\mu_{21} + \mu_{03})[3(\mu_{30} + \mu_{12})^2 - (\mu_{21} + \mu_{03})^2] \\
I_6 &= (\mu_{20} - \mu_{02})[(\mu_{30} + \mu_{12})^2 - (\mu_{21} + \mu_{03})^2] \\
&\quad + 4\mu_{11}(\mu_{30} + \mu_{12})(\mu_{21} + \mu_{03}) \\
I_7 &= (3\mu_{31} - \mu_{03})(\mu_{30} + \mu_{12})[(\mu_{30} + \mu_{12})^2 - 3(\mu_{21} + \mu_{03})^2] \\
&\quad - (\mu_{30} - 3\mu_{12})(\mu_{21} + \mu_{03})[3(\mu_{30} + \mu_{12})^2 - (\mu_{21} + \mu_{03})^2]
\end{aligned} \tag{3.15}$$

The primary benefits of using invariants like this is that they do not require any manipulation of the original image, and, when calculated using scale and translationally invariant values, they are invariant to translation, rotation, and scaling. I_7 is also invariant to mirroring of the image.

Since their discovery, Hu's seven invariants have shown their usefulness in a number of

different applications, including aircraft [98], palmprint [99], and plant leaf [100] identification. However, Flusser [101] has shown that these seven invariants are not independent or complete. He achieved this by first detailing a method of generating rotationally invariant moments of any order which are based around complex moments: moments created by using a complex basis function. It was then shown that Hu's invariants were combinations of these more basic moments. Flusser's moments will be used later in this work.

Based on what is called the *generalized fundamental theorem of moment invariants* (GFTMI) [102], Mamistvalov constructs a number of different moment invariants based on affine transformations and their subgroups [103]. Affine transformations are those which preserve straight lines. These invariants which were constructed include a parameter shown in equation 3.16. MacSleyne *et al.* [104] found that combining one of Hu's moments with this value of Mamistvalov's results in two quantifiers which are especially suited to quantifying 2D shapes and are used extensively in this work:

$$\Omega_1 \equiv \frac{A^2}{\mu_{20} + \mu_{02}} \tag{3.16a}$$

$$\Omega_2 \equiv \frac{A^4}{\mu_{20}\mu_{02} - \mu_{11}^2} \tag{3.16b}$$

It was noted that Ω_1 is representative of the aspect ratio of the shape, with larger values indicating a ratio closer to one, while Ω_2 is indicative of the complexity of the shape. Shapes with a smaller perimeter to area ratio will have larger values. For both of these values the maximum is generated by a circle. The implications and use of these values with special attention to their use in superalloy precipitate shape measurement will be discussed in Chapter VI.

Callahan *et al.* also noted that, when plotted against each other, these values generate a *second order moment invariant map* (SOMIM) [9], which is shown in Figure 3.4. Because ω_1 is sensitive to affine transformations, shapes which generate points which are on the same horizontal line on the SOMIM differ in aspect ratio, but are of the same class. ω_2 is invariant to affine transformations, so as one moves vertically top to bottom, the complexity of the shapes increases. Below the line indicating triangles, the classes of shapes becomes less identifiable from the ω_2 value alone. Shapes of note are marked in the figure, and the

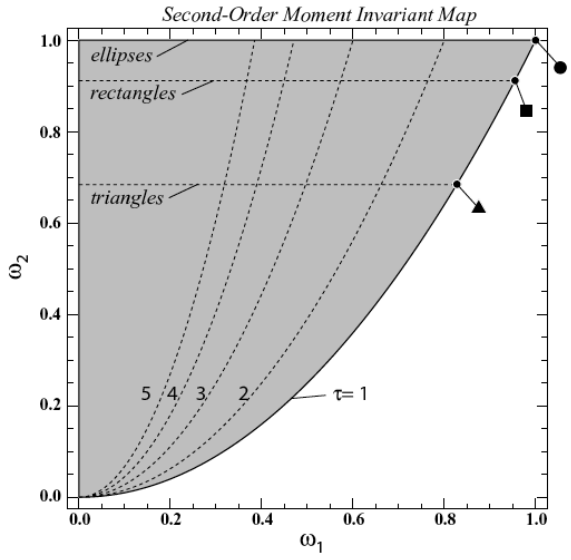

Figure 3.4: Second order moment invariant map (SOMIM) generated from plotting ω_1 against ω_2 normalized by the values of the cirlce. From [9].

gray region indicates possible locations for shapes. No shapes exist which would generate points in this space outside of this region.

These moments only utilize moments up to the second order. It can be shown that the moments up to the second order only allow reconstruction of a shape to its equivalent ellipse. Any information not contained in this ellipse is lost. This includes information like area, aspect ratio, and relative perimeter length. More complex shape information can be captured using moments of a higher order.

Similar to the ω parameters, but of higher order, are those of τ_1, τ_2, and τ_3, which were first described by Flusser and Suk [105] and modified and adjusted by Callahan *et al.* [9] for use in 2D γ' precipitate analysis. These values are as follows and based upon moments up to the 4th order, with normalizations to set the maximum value to 1:

$$\tau_1 \equiv \frac{\mu_{00}^3}{3\pi^2}(\mu_{40} + 2\mu_{22} + \mu_{04})^{-1} \tag{3.17a}$$

$$\tau_2 \equiv \frac{\mu_{00}^6}{48\pi^4}(\mu_{40}\mu_{04} - \mu_{31}\mu_{13} + 3\mu_{22}^2)^{-1} \tag{3.17b}$$

$$\tau_3 \equiv \frac{\mu_{00}^9}{1728\pi^6}(\mu_{40}\mu_{22}\mu_{04} - \mu_{40}\mu_{13}^2 - \mu_{04}\mu_{31}^2 + 2\mu_{31}\mu_{13}\mu_{22} - \mu_{22}^3)^{-1} \tag{3.17c}$$

Utilizing these values gives a more detailed picture of shape changes near the exterior of

the shapes. In addition, both τ_2 and τ_3 are affine invariants. Once again, the maximum for each of these values is generated by the circle.

Beyond the fourth order, Suk and Flusser [106] have developed a method for calculating affine moment invariants of arbitrary order, which allow for even further detail to be quantified. However, with increasing order comes the requirement for increased computational resources. In addition, higher order moments are more sensitive to noise, which must be balanced against the increased detail that can be achieved through these descriptors.

3.1.3 Three-dimensional Cartesian moments

As discussed in Section 2.3, three dimensional shape analysis provides a more complete view of the nature of an object, but can drastically increase the complexity of shape measurement and descriptor calculation. Fortunately 3D Cartesian moments are not much more complicated than their 2D counterparts, but typically require more resources for calculation.

Methods for calculation and the forcing of translation invariance are very similar to those for 2D moments. This includes the general formula for 3D Cartesian moments:

$$m_{pqr} = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} x^p y^q z^r f(x,y,z)dxdydz \tag{3.18}$$

and the calculation of central moments:

$$m_{pqr} = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} (x-\bar{x})^p (y-\bar{y})^q (z-\bar{z})^r f(x,y,z)dxdydz \tag{3.19}$$

where $\bar{x} = \frac{m_{100}}{m_{000}}$, $\bar{y} = \frac{m_{010}}{m_{000}}$, $\bar{z} = \frac{m_{001}}{m_{000}}$.

While rotational invariance can be achieved in a somewhat similar way to the principal axis method which was suggested by Hu for 2D, this is again undesirable, as the rotation of the 3D object can lead to distortion at the boundaries, which play a larger and larger role in the values of the moments of increasing order.

Invariants

Utilizing moments of the second order, the *moment of inertia tensor*, I, can be generated:

$$I = \begin{pmatrix} \mu_{020} + \mu_{002} & -\mu_{110} & -\mu_{101} \\ -\mu_{110} & \mu_{200} + \mu_{002} & -\mu_{011} \\ -\mu_{101} & -\mu_{011} & \mu_{200} + \mu_{020} \end{pmatrix} \tag{3.20}$$

Utilizing the invariants of this tensor, and normalizing by the volume of the object for scale independence, MacSleyne *et al.* identified the following three invariants:

$$\Omega_1 \equiv 3V^{5/3}(\mu_{200} + \mu_{020} + \mu_{002})^{-1} \tag{3.21a}$$

$$\Omega_2 \equiv 3V^{10/3}(\mu_{200}\mu_{020} + \mu_{200}\mu_{002} + \mu_{020}\mu_{002} - \mu_{110}^2 - \mu_{101}^2 - \mu_{011}^2)^{-1} \tag{3.21b}$$

$$\Omega_3 \equiv V^5(\mu_{200}\mu_{020}\mu_{002} + 2\mu_{110}\mu_{101}\mu_{011} - \mu_{200}\mu_{011^2} - \mu_{020}\mu_{101}^2 - \mu_{002}\mu_{110}^2)^{-1} \tag{3.21c}$$

MacSleyne *et al.* also noted that the maximum value for these invariants is generated by the sphere. These values generate a 3D space which is not dissimilar to the SOMIM shown in Figure 3.4. Shapes are restricted to a smaller region within what, in this case, is a unit cube, with regular shapes restricted to a curve through the space. Boundaries of this region represent certain morphological extremes.

Suk and Flusser [107] describes a generalization of the 2D method used in [106] that is used to generate 3D moment invariants up to the 16^{th} order. Although not all of these are independent, they provide a way to capture a large degree of shape information in a manner that is invariant to basic transformations which allow improved object recognition.

3.2 Zernike moments and invariants

3.2.1 Two-dimensional Zernike moments

In order to address the issue of redundancy (non-orthogonality) in the different orders of Cartesian moments, Khotanzad and Hong [95] introduced the idea of using Zernike polynomials as basis functions for moment generation. These moments would share the properties of the polynomials, which were formulated for their orthogonality over the unit

circle. Zernike polynomials have the form:

$$V_{nm}(x,y) = V_{nm}(\rho,\theta) = R_{nm}(\rho)e^{im\theta} \tag{3.22}$$

where $n \geq 0$ and is an integer , m is all integers which satisfy $(n-|m|)$ is even and $|m| \leq n$, $\rho = \sqrt{x^2+y^2}$, and θ is the angle in radians from the vector (x,y) to the x-axis, in the counterclockwise direction.

$R_{nm}(\rho)$ is the radial component of this polynomial, as it only contains information regarding the radial distance to the point (x,y). This function is defined as:

$$R_{nm}(\rho) = \sum_{s=0}^{\frac{n-|m|}{2}} (-1)^s \frac{(n-s)!}{s!(\frac{n+|m|}{2}-s)!(\frac{n-|m|}{2}-s)!}\rho^{n-2s} \tag{3.23}$$

Using these polynomials, Zernike moments take the form:

$$Z_{nm} = \frac{n+1}{\pi}\iint_C V^*_{nm}(\rho,\theta)f(x,y)dxdy \tag{3.24}$$

where C is the unit circle. It should also be noted that $Z^*_{nm} = Z_{n,-m}$. Because of this, it is sometimes specified that $m \geq 0$. This is because for shape description, it is only the magnitude of the moments which is relevant.

In analysis of five separate types of 2D moments, including Cartesian, complex, Zernike, pseudo-Zernike, and Legendre moments, Teh and Chin noted that the Cartesian and complex, which do not have orthogonal bases, performed the worst in image reconstruction, requiring more orders to achieve the same quality of image. They also noted that Zernike and pseudo-Zernike performed best overall in terms of showing the least sensitivity to noise and in reconstruction [108].

Comparing the radial portion of the Zernike basis function (Figure 3.5) to that of the Cartesian moments (Figure 3.2) illustrates an advantage of Zernike moments which contributes to their improved image representation as well. The difference between functions of increasing order is more pronounced for Zernike moments and the values for some of these functions at coordinates closer to zero is larger, resulting in more information about the

center of the object. This is less important for regular shapes, but improves shape representation for more irregular shapes with boundaries that may lay near the centroid. Other moments, such as orthogonal Fourier-Mellin moments (OFFMs) have been proposed based on improved coverage of this region and general behavior of the basis functions [109].

Because Zernike moments are restricted to the unit circle, the shape in question must be scaled to fit within the circle, and translated to the origin as well. This necessitates the calculation of Cartesian moments for the shape prior to the calculation of Zernike moments.

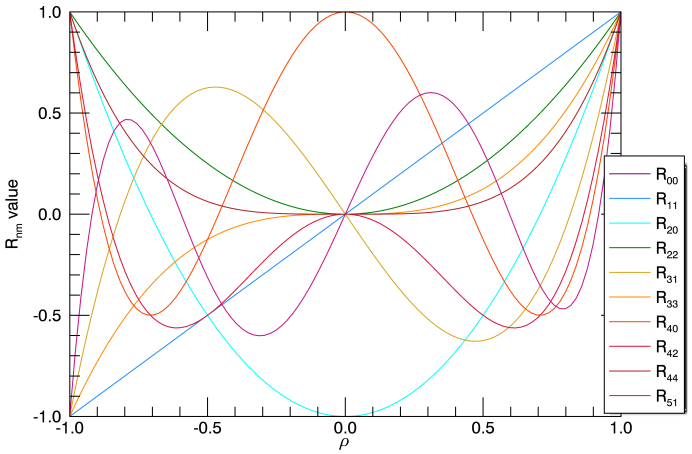

Figure 3.5: Radial function (R_{nm}) for first ten Zernike polynomials.

Invariants

Khotanzad and Hong [95] also noted that, for some object rotation, α, with a given Zernike moment Z_{nm}, the resulting Zernike moment is given as $Z_{nm}e^{-jm\alpha}$. Therefore the magnitude of a given moment, $|Z_{nm}|$ is rotationally invariant. When computed for five separate rotations, it was found that the resulting magnitudes of the moments only varied by a maximum of 0.9% for Z_{22}, which was attributed to the discrete nature of the image and distortion at the edges due to changes in pixels.

Belkasim *et al.* established a method of selecting Zernike moments as invariant descriptors that prevents the use of dependent values. A normalization scheme is also suggested, as well as a method of calculating these normalized Zernike moments from central Cartesian moments. This provides an advantage because it avoids any discrepancies which can be brought about from the translation and scaling required for initial Zernike moment calculation.

3.2.2 Three-dimensional Zernike moments and invariants

Novotni and Klein [10], working with the 3D Zernike moments first established by Canterakis [110], compare the properties of these descriptors with others in terms of shape retrieval and calculation resource intensity. The derivation is complex and will not be described here, however it are based upon spherical harmonics and Legendre functions. The final form of the 3D Zernike moments are:

$$Z_n^m l(r) = \sum_{\nu=0}^{\frac{n-l}{2}} \sum_{\kappa}^{\nu} \sum_{\alpha=0}^{m} \sum_{\mu=0}^{|\frac{l-m}{2}|} \sum_{\beta=0}^{\kappa+\mu} \frac{(-1)^{m+\mu}}{2^{m+2\mu}} q_{nl}^{\nu} c_l^m i^{m-\alpha} \binom{\nu}{\kappa} \binom{m}{\alpha} \binom{\kappa+\alpha}{\beta} \binom{l}{\mu} \binom{l-\mu}{m+\mu} \times x^{2\beta+\alpha} y^{2(\kappa+\mu-\beta)+m-\alpha} z^{2(\nu-\kappa-\mu)+l-m} \tag{3.25}$$

Novotni and Klein conducted reconstruction using these descriptors and found that by the 20^{th} order, the shape was reconstructed to a good level of fidelity. This is shown in Figure 3.6. They also found that 3D Zernike descriptors outperformed similar descriptors based entirely upon spherical harmonics for object retrieval.

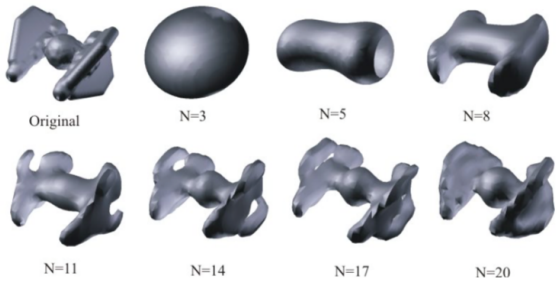

Figure 3.6: Reconstruction of of a spaceship model using 3D Zernike moments. From [10].

CHAPTER IV

Cell function identification in *Arabidopsis Thaliana* cotyledon through 3D shape analysis

4.1 Background

The British metallurgist, scientific historian, and Manhattan Project researcher Cyril Stanley Smith wrote in his essay "The shape of things" in volume 190, issue 1 of the *Scientific American*: "Statistical studies...have shown that the shapes of soap bubbles match closely those of vegetable cells, metal grains and animal cells." He goes on to say, "In sum, the structure of any aggregate of matter - whether organic or inorganic, biological, natural or artificial - depends not only on the interplay of forces but also on the very simple but inescapable mathematical requirements of space-filling" [111]. In this work, and others like it [112], Smith remarked on how similar shapes reappear again and again in nature.

This work attempts to analyze how three-dimensional shape can be used to identify the function of cells in the embryonic leaves, or cotyledons, of the plant *Arabidopsis Thaliana*. By analyzing feature sets and establishing methods by which form and function can be linked in a system with clearly delineated classes such as these leaves, it is hoped that, as Smith suggests, what is learned here can be extended to animal cells, or metal grains.

4.1.1 Shape analysis in cell and tissue analysis

Shape is powerful tool for the analysis of cells. It is well known that cell morphology and function are linked. According to Pincus and Theriot [113], who conducted a survey and comparison of 2D cell morphology descriptors, "Cell shape is a large-scale expression of many organizing, competing and highly regulated biological processes ..." To use the old

adage, *form follows function*. By quantifying cell shape, the investigator can hope to gain insight into the function of the cell through its form.

Because of this relationship between these two aspects of cells, shape has become a commonly included feature in image cytometry, a fundamental method of cell characterization. Over the past fifty years, an extensive amount of work has been focused on automating the cytometry process, driven by the obvious benefits of computer-based analysis: speed and precision. These characteristics are particularly important when dealing with the large numbers of cells that are usually involved in such an analysis. Automated analysis also provides solutions to a number of potential problems with a human-based analysis, primarily: reduced quantity and quality of gathered information and speed of analysis.

In general, automated analysis can be preferable to manual analysis, especially in situations where large amounts of data can be made available for processing by computer systems. However, the most basic kind of data in these fields, images, does not lend itself to automated analysis. While a scientist may be able to visually identify regions of interest and perform measurements on them, creating an algorithm or a set of rules by which to do so automatically is much more difficult. Human judgment cannot easily be coded into a program. Given the prevalence of image-based data and the desire to take advantage of ever more capable computer systems, significant amounts of work have focused on improving the performance of automated image analysis. In the last decade, machine learning techniques have dramatically improved the quality of segmentations and allowed for better identification of meaningful features.

Automated analysis in the context of biological materials can be undertaken in two or three dimensions, although 2D images are by far the most common. 2D images are favored because of the difficulty in capturing 3D data for softer biological materials, and because an entire class of images, that of cells on microscope slides, only exists in two dimensions. Typical features examined include 2D shape descriptors such as aspect ratio or circularity, as well as global descriptors like moment invariants or boundary descriptors like Fourier descriptors (discussed in Chapter II). While the materials scientist usually works with electron micrographs which lack any variability in color, this is not true in this context, so color and texture are commonly utilized as well. In three-dimensional analysis, the typical

descriptor is a global shape descriptor, although once again this kind of analysis is typically difficult to conduct due to the nature of the samples.

Shape provides a window into these interactions and allows for the differentiation of cell based on their function. To this end, Pincus and Theriot analyzed three representative cell shapes, those of a bacterium, kidney cells, and keratocyte cells. Shape descriptors examined were Fourier descriptors and Zernike moments, although raw distance maps, binarized images, and outlines were utilized also. Principal component analysis was also used as a method of improving reconstructions by maximizing the value of the first several parameters in a reconstruction. It was found that signed distance maps, combined with PCA analysis to pull out the most relevant aspects of the data, perform the best.

In 1966, Prewitt and Mendelsohn [13] conducted one of the first experiments on automated cell classification. Four of the most visually distinct white blood cell types were examined. These were neutrophils, eosinophils, lymphocytes, and monocytes. These cells were extracted from human blood and stained. Then, by rastering light through the eyepiece of an optical microscope, optical density measurements were taken every 0.25 microns, which, when scaled over the range of 0-255, took the form of 200×200 pixel images. A large number of features were calculated for each cell, and then the range of features analyzed for maximum power of discrimination. It was found that a three dimensional space formed by the area of the nucleus, the area of the cytoplasm, and the contrast a the outer diameter allows for sufficient discriminatory power to automatically identify these cells. It was noted that shape analysis would be a useful additional tool.

Similar work in white blood cell analysis was undertaken by Ongun *et al.* [114], who utilized features based upon basic shape information, affine Cartesian moment invariants, CIE L*a*b* color space, and texture of the images to generate a total of 57 features from which to draw upon. Using four different classification methods and by treating each cell as a single 57 dimensional vector, it was found that support vector machines (SVMs) were the most effective method with 91% accuracy on the test set.

Aside from blood cells, shape descriptors have been used to conduct automated identification of breast cancer tissue. Tahmasbi *et al.* [115] set out to improve computer-assisted diagnosis (CAD) systems through the application of 2D Zernike moments to a set of mam-

mograms with unidentified masses present. For each mass, Zernike moments from order 3 to 17 were calculated for both the entire mass and the margin alone, resulting in 64 individual moments, and 128 features overall. Using a multilayer perceptron (MLP) as a classifier, the authors were able to achieve an accuracy of 96% while maintaining a false negative rate of 0%. This was particularly important, because a false negative in medicine is particularly harmful. Segmentation in this case was manually performed, so it cannot be considered a truly automated process.

Sharma and Khanna [116] attempted to further improve the ability of a CAD system to identify cancerous masses in mammograms while also fully automating the process. To this end, each mammogram is preprocessed by first applying histogram equalization to normalize the contrast across the image. The pixels of the mammogram are then partitioned to three grey-levels using the *k-means* algorithm, which effectively performs a tri-level segmentation of the image with dynamic threshold values for each image. This separates pixels into low, medium and high density values. Patches of the mammogram are then extracted and the entropy calculated for each. Regions with large entropy have the most texture information, and these are further tested for malignancies. This testing is performed using Zernike moments up to order 20 with SVMs as the classifier. The result showed 99% sensitivity and 99% specificity.

Three dimensional cell analysis is relatively rare. Although observations regarding 3D morphologies, and their change with regard to a number of different interior and exterior factors are noted, there are very few studies which involve the quantification of three dimensional cell shapes. Khairy *et al.* [117] utilized spherical harmonics as basis functions for the quantification of blood cell morphology and found that the membrane cytoskeleton "can account quantitatively for a wide range of RBC [red blood cell] morphologies..." Elliott utilized [118] a combination of graph based methods, surface curvature analysis, and the formulation of an invariant intra-cell coordinate system to measure three dimensional shape of cells undergoing migration and to correlate them to fluorescence signals originating within the cell.

The lack of focus on three-dimensional information is a direct result of the difficulties in gathering this type of information. However, recent advances in cell data acquisition are

making 3D image data easier to collect [119,120]. This is important because cells are three-dimensional objects which occupy and operate within a three-dimensional environment. Establishing features and methods for the analysis of cell shape in this context is important for the future of cellular analysis.

This work attempts to answer some of these questions by focusing on the use of Cartesian and Zernike moment invariants as descriptors of cell shape. By exploring these features in three dimensions, we hope to establish the extent to which shape can be used to differentiate healthy *Arabidopsis Thaliana* cotyledon cells (epidermal, spongy mesophyll, pallisade mesophyll, and vascular) from each other. Distinguishing healthy cells from a system with distinct categories allows for the better establishment of a ground truth and provides a stepping stone to the identification of irregular cells, the development of a general classifier for an entire species, and eventually the application of this method to other systems.

4.1.2 Arabidopsis Thaliana

The plant *Arabidopsis Thaliana*, also known as thale cress or simply arabidopsis, is a small flowering plant which is considered a weed.The first mutant of this species was recorded in 1873, and in 1943 its use as a model organism was proposed. It was the first plant to have its genome sequenced in 2000 by the Arabidopsis Genome Initiative [121], and has one of the smallest genomes among plants.

The data analyzed in this study are image stacks from two *A. Thaliana* cotyledons. Cotyledons are the embryonic leaves which exist within the seed. Upon germination, they form the first visible leaves of the plant. Cotyledons consist of five major types of cells. Epidermal cells make up the outer layer and form a protective barrier, they are more irregular in shape and are most visible in Figure 4.1c and 4.2a. There are mesophyll cells of two types, palisade and spongy, which comprise the majority of the tissue and act as the primary photosynthetic cells in the plant. Palisade cells are more cylindrical in shape and appear along the edges of the leaf in Figures 4.1a, 4.1b, and 4.2b. Spongy mesophyll are more spherical and make up the majority of the cells. They are especially prevalent in Figures 4.1a, 4.1b, 4.2b, and 4.2c. Stoma facilitate gas exchange, tend to be more spherical, are visible in between epidermal cells. Finally, proto-vascular or protophloem cells allow for

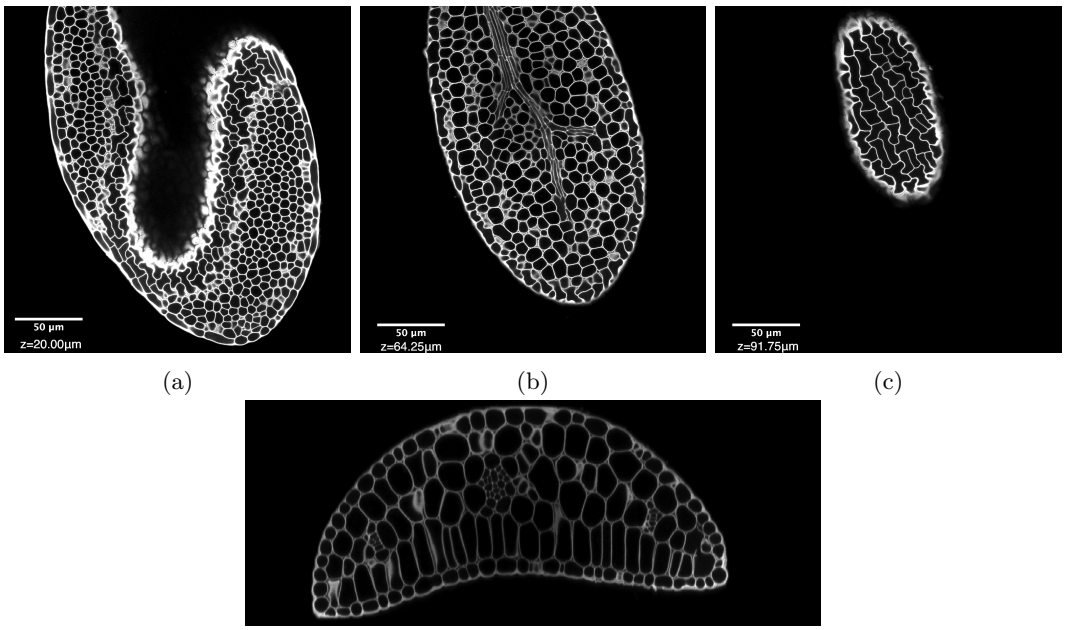

(a) (b) (c)

(d)

Figure 4.1: Examples of 7x series images. (a-c) example images from along the z-axis. (d) shows a image taken along the x-axis with more cylindrical palisade mesophyll visible along the bottom of the cotyledon.

the movement of liquid into to and out of the plant, carrying water and sugar. These are the distinctly long cells in Figures 4.1b and 4.2b.

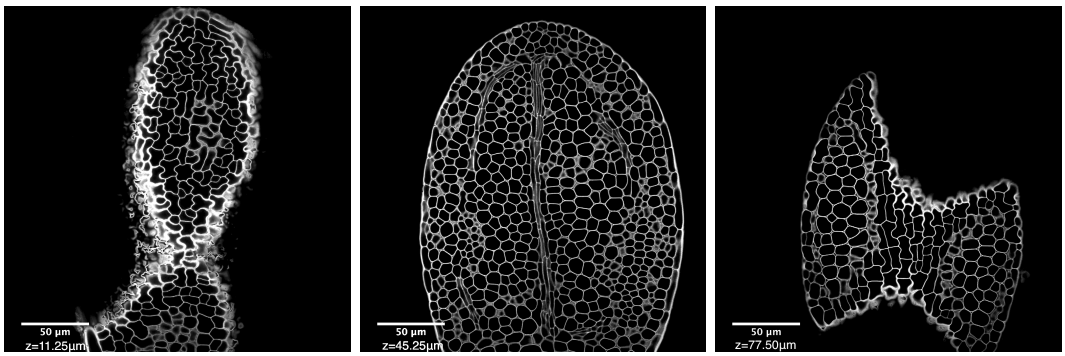

(a) (b) (c)

Figure 4.2: Examples of 8x series images.

4.2 High dimensional analysis

Large feature vectors, like the kind used here, necessitate a method of analysis that is more adapted to handling the problems that occur in these unintuitive, high dimensional spaces. Bellman coined the phrase "curse of dimensionality" [122] to refer to the problems that arise when operating with data in high dimensional spaces that do not occur in more intuitive, lower dimensional spaces. One of the central problems, and the one that is most troublesome when attempting to determine similarity or dissimilarity of objects in shape space, is the loss of meaning that occurs for the Euclidean distance function as the number of dimensions grows.

To illustrate this point, Zimek *et al.* [123] describe how the variance of the ratio length of any point vector ($||X_d||$) to the mean point vector $E[||X_d||]$ converges to zero with increasing dimension, which implies the following:

$$\text{If } \lim_{d\to\infty} \text{var}\Big(\frac{||X_d||}{E[||X_d||]}\Big) = 0, \qquad \text{then} \quad \frac{D_{\max} - D_{\min}}{D_{\min}} \to 0 \tag{4.1}$$

where $D_{\max}$ is the maximum distance between two points of a set, and $D_{\min}$ is the minimum distance. This applies to any L_p norm with $p \geq 1$. So, picturing a set of d-dimensional points as they are embedded in higher and higher dimensional spaces, the L_p norm between any two points will decrease with increasing n, even if the points themselves do not change relative to each other. Essentially, as the number of dimensions increases, the distances between any two points become more and more similar.

Because the shape-data generated through calculating moment invariants of a range of orders is high-dimensional, it is necessary to employ analysis techniques which can help address the issues inherent in these types of spaces. A tremendous amount of work has been done in the field of machine learning to develop methods for these situations. A primary application of machine learning is the assignment of class labels to individual points of data. Two different techniques, considered *supervised* and *unsupervised* respectively, were chosen and are detailed in this section.

4.2.1 Supervised machine learning

Supervised machine learning is a broad field, but with respect to data classification can be described as the generation and application of a model for the assignment of class labels for a given data set based on an example, or training, set which has be pre-classified through some other process (although usually done by hand by an observer). There are a wide variety of supervised machine learning algorithms available, each with a variety of advantages and disadvantages. These include decision trees [124], support vector machines [125], and neural networks [126]. For an overview of the relative strengths of each method, the reader is directed to Kotsiantis [127]. Support vector machines were utilized in this work because of their simplicity, relative accuracy, and ease of implementation.

Support vector machines

Support vector machines (SVMs) were first developed by Vapnik and Chervonenkis in 1964 and detailed in their book [128]. They addressed the problem of identifying the optimal separating hyperplane for two *separable* classes of data in an p-dimensional space where these classes are known. Given a set of data:

$$(y_1, x_1), \ldots, (y_n, x_n) \tag{4.2}$$

where x_n represents a point in a p-dimensional space, and y_n its class with $y \in \{-1, 1\}$. Two parallel hyperplanes can be constructed, each touching one class of the data as:

$$\begin{aligned} \boldsymbol{w} \cdot \boldsymbol{x} - b &= 1 \\ \boldsymbol{w} \cdot \boldsymbol{x} - b &= -1 \end{aligned} \tag{4.3}$$

The distance between these two planes:

$$d = \frac{2}{\| \boldsymbol{w} \|^2} \tag{4.4}$$

is maximized when $\| \boldsymbol{w} \|^2$ is minimized. The single optimal hyperplane will therefore exist

at

$$\boldsymbol{w} \cdot \boldsymbol{x} - b = 0 \tag{4.5}$$

with the minimal $\| \boldsymbol{w} \|^2$ subject to the constraint:

$$y_i(\boldsymbol{w} \cdot \boldsymbol{x}_i + b) \geq 1, i = 1, \ldots, n \tag{4.6}$$

which ensures that no points are misclassified. This system is then optimized using Lagrange multipliers. A full derivation and discussion of the optimization can be found in [125].

SVMs ultimately reach their modern form in 1995 with Cortes and Vapnik [129]. Here, the authors did *not* assume that the two classes were separable, but instead implemented a cost function which was equal to zero in the case of correctly classified points but otherwise proportional to the distance of a misclassified point from the given hyperplane. Identifying the optimal hyperplane is then a matter of minimizing equation 4.4 combined with the value of the cost function.

SVMs have proven to be one of the best-performing supervised learning methods, nearly as good as neural networks in many applications and in some cases surpass them [130]. They also have the benefit of being much faster to implement and easier to understand.

4.2.2 Unsupervised machine learning

In contrast to supervised methods, unsupervised machine learning does not rely on pre-classified training sets in order to develop models. Instead, the algorithms are constructed to find underlying characteristics of the data without any input from the user (with the exception of parameters). This provides a number of benefits, the first and most obvious of which is the user need not construct a training set. This saves time in the initial setup of the algorithm and also prevents the user from inadvertently influencing the results. This can be beneficial in that it can prevent errors from improper training or biases of the user which can be passed to the model. However, without user input, the model may not generate expected results. The unsupervised machine learning method described below was utilized to determine what could be learned regarding the nature of the cells from shape and the other feature values utilized alone.

In this work, the unsupervised method was performed in two steps. The first is a dimensionality reduction method, which serves to map the high-dimensional data into a more workable space while attempting to retain the most important spatial aspects of the data. This is done by either finding directions in the original space which better describe the data than the original axes, or by mapping data points in high-dimensional space to positions in a lower-dimensional space. The next step is a clustering technique, which groups points that are nearby each other in this new space, with the underlying assumption that proximity indicates similarity. These methods are described in detail below.

Dimensionality reduction

The first method used to achieve a reduction in dimensionality is principal component analysis (PCA). This method is relatively simple, but allows for reduction to an arbitrary degree and provides a good first step in any kind of dimensionality reduction scheme and is the first step in the technique discussed here. PCA works by identifying the n orthogonal directions (or principal components) within an d-dimensional data-space which account for the highest variability within the data. These directions are ranked according to the amount of variability accounted for by each. The data is then projected onto these axes, and by selecting the first x directions, allows for the projection of the d dimensional space to one which can be visualized while preserving as much variability as possible. This implies that the direction with the highest variability is inherently the most informative, and that correlations between dimensions are linear.

Practically speaking, PCA is performed on an m dimensional data set with n number of points ($m \times n$). First, the data is translated so that the mean of each column is coincident with the origin. This facilitates the calculation of the correlation or covariance matrices, which captures the relationship between dimensions of the data and gives a rough idea of the orientation of the point cloud generated in data-space. Traditional PCA uses the covariance matrix, but in most cases, the correlation matrix is utilized for the calculation of the principal components in order to standardize the values. The eigenvectors are then calculated for the resulting $m \times m$ matrix and sorted by decreasing eigenvalue. With m eigenvalues (λ_i) the fraction of the variability accounted for by the i^{th} eigenvector is:

$\frac{\lambda_i}{\lambda_1+\lambda_2...+\lambda_m}$.

PCA is a simple linear projection, and can miss some of the more complex structures the data can form in a high-dimensional space. In order to address these issues and provide a method which preserves this information through dimensionality reduction, Maaten and Hinton [131] developed a method known as t-distributed stochastic neighbor embedding or t-SNE. This method is based upon stochastic neighbor embedding developed by Hinton and Roweis [132]. This method necessarily has to be nonlinear in order to capture the nature of the data distribution.

Once again for an $m \times n$ dataset, t-SNE works by fitting an m-dimensional Gaussian with a fixed variance to each data point. The value for point j of the Gaussian centered at point i is the conditional probability ($p_{j|i}$) that j is similar to i. Using this method an $n \times n$ similarity matrix can be constructed containing this probability for every point. The use of a Gaussian implies that the value will exist for every point but will be very small for points far away. The same thing can be done in a low dimensional space with the same data, giving a matrix $q_{j|i}$, except in this case, the distribution used is not a Gaussian but a t-distribution with heavy tails to reduce crowding in the final map.

The goal is then to distribute the points in the low dimensional space in such a way that $p_{j|i}$ is as close to $q_{j|i}$ as possible. The cost function used to minimize the differences between these distributions is the Kullback-Leibler divergence. This function is minimized through gradient descent to render the final map. The primary parameter that must be passed into this is the *perplexity*, which controls the spreads of the probability matching functions and effectively controls the degree of influence that each point exerts on all other points. t-SNE combined with machine learning has been used extensively in a wide variety of fields over the last decade [133–137].

Feature clustering

Once the application of t-SNE has generated a map, the resulting points in two-dimensions can then be clustered using proximity on the resulting 2D map to indicate shape similarity. t-SNE is a powerful tool that preserves the structure of the high-dimensional data, but that structure is not always as expected. Using traditional clustering methods such as k-means

clustering assumes Gaussian or globular cluster shapes which is not always the case. In addition k-means is more of a partitioning algorithm, rather than a clustering technique, the difference being that k-means partitions the entire space into k number of clusters, but is incapable of identifying outliers. Because the final goal of this work is to identify cell function through morphology, it would be desirable if irregular or underrepresented cells, such as malformed, mutated, or cells mid-mitosis could be automatically identified. These irregular shapes should not belong to any cluster, so a method is needed that can pick them out.

Hierarchical density-based spatial clustering of applications with noise, or HDBSCAN, is such a technique [138]. HDBSCAN works by utilizing *mutual reachability distance* rather than Euclidean distance. This is defined as: $d_{\text{mreach, k}}(a,b) = max\{\text{core}_k(a), \text{core}_k(b), d(a,b)\}$, where $\text{core}_k(a)$ is the radius of a circle centered on point a that contains k other points, or alternatively, the distance to the k^{th} nearest neighbor. $d_{\text{mreach, k}}$ is calculated for every point and essentially moves outlying points even farther away, accentuating the closeness of the other points.

Next, a *minimum spanning tree* is created by adding points to the tree, always adding the shortest length edge that connects a non-member point to the tree. The edges are then ranked in order of length to create a hierarchy. This is iterated through, shortest to largest, and each edge and the clusters it links together are compiled. Finally, utilizing a user-provided parameter: *minimum cluster size*, the edges combining clusters are iterated through, longest to shortest, and only those which connect clusters greater than this parameter are retained. The final list of edges is then searched for those which existed longest and these are selected as final clusters. This method is powerful and provides noise identification, which is ideal for the purposes of this work. A tutorial on HDBSCAN can be found at [19].

4.3 Materials and methods

4.3.1 *Arabidopsis Thaliana* cotyledons

Two sets of cross-sectional images, one from each of the two samples (henceforth referred to as the 7x series and 8x series) were provided by Dr. J.C. Palauqui of the National

Institute of Agricultural Research (INRA), Paris, France. Details regarding the method of the collection of this data are given in [120] and are summarized here. The primary imaging method for these two samples is confocal laser scanning microscopy (CLSM). This technique is popular in the analysis of biological specimens for a number of reasons. CLSM requires the relevant tissue to be fluorescently labeled. That is, the regions to be imaged must first be stained with a dye that fluoresces under a given wavelength of light. Identifying a proper dye is an entirely different problem which can be difficult to solve in its own right. Truernit *et al.* utilized a modified pseudo-Schiff propidium iodide staining technique (mPS-PI), which provides superior staining ability to previously employed methods and allows for better resolution when combined with CLSM.

The primary benefit of CLSM is it allows the imaging of material within the bulk of the sample. This is due to the fluorescence of the material being imaged. Problematically, the cotyledons being observed here have evolved to reduce the propagation of radiation into the bulk material, which is to be imaged. However, the mPS-PI technique implements a clearing solution which makes these regions less opaque. Not only is CLSM capable of imaging interior regions, but these may be selectively viewed through a process called *optical sectioning*. CLSM works by rastering a laser beam across the sample. The depth from which fluorescent light is captured can be selected by using a spatial filter, such as a slit or pinhole. This blocks out unwanted fluorescence from other z levels and from general fluorescence of the specimen. This allows very precise visualization of specific levels within the sample, without any mechanical sectioning whatsoever. Vertical resolutions of $0.23\mu m$ can be achieved under optimal conditions [139].

4.3.2 Images and image processing

Three sets of representative example images are shown in Figure 4.2 along with the indicated z level. Each one represents an upper, lower, and middle region, and in each, examples of the five different cell types are indicated. Sample set 7x has dimensions $256 \times 256 \times 100\mu m^3$, while sample set 8x has dimensions $256 \times 256 \times 96.25\mu m^3$.

Individual images were segmented using *ilastik*, described in Section 6.2.2. This program utilizes decision trees to determine to which class each pixel of the image belongs. Using

ilastik results in segmentations with an average out of bag error of less than 1%. This error is, and can only be, based upon the training set. Estimating total segmentation effectiveness would require a ground truth which does not exist.

After segmentation, each two-dimensional slice was opened (image erosion followed by dilation) by a small, 3×3 kernel in order to improve final segmentation by removing small pixels connecting two regions which should be separate. Despite the quality of the initial segmentation, this is necessary due to inconsistencies in cell wall contrast and because even a single connecting pixel between two regions is enough to merge two regions and add error to the analysis. Shapes with a volume of less than 50 voxels were excluded from the analysis in an effort to avoid introducing errors into the t-SNE mapping.

Similar to the methods described in 6.2.2, Interactive Data Language, IDL, was used to calculate the shape descriptors used in this analysis. First, each shape was labeled with an individual value. Then the shape was extracted from the bulk and analyzed. The shape quantities were calculated using the method of Novotni and Klein [10], where first the central moments were calculated for each shape, and normalized according to the volume of the object. Finally, they are converted into the relevant shape values, as described below.

4.3.3 Ground truth and scoring

A ground truth was established by simple manual identification of cells chosen at random from the set of all cells. These cells were either labeled as epidermal, vascular, spongy mesophyll, and palisade mesophyll. For both the 7x and 8x cotyledons, the size of the truth set was 1251 cells, roughly equal to 1/3 of the size of the entire set (3585 for 7x and 3712 for 8x).

Calculating the success of an unsupervised classification method requires a metric that is capable of quantifying the similarity of two sets that satisfies two conditions: the metric must be successful when labels do not match and the metric must be interpretable. The first point is particularly important in the case of unsupervised machine learning, where labels may conform to the known class labels but may be of different values (*i.e.* all points of class 1 being assigned a value of 2). The metric must also handle two problematic cases: a classification that gives a different label to each cell and a classification that assigns a

single class to the entire set. In the case of the first, a metric which measures the purity of each assigned class will give a false perfect result, when obviously this is not useful. In the second case, a purity metric on the truth labels fails, and can return a perfect result.

One metric which satisfies these requirements is the *normalized mutual information* (NMI) shown in equation 4.7.

$$\mathrm{NMI}(X;Y) = \frac{2I(X;Y)}{H(X) + H(Y)} \tag{4.7}$$

where $I(X;Y)$ is the mutual information of set X given Y, and $H(X)$ is the entropy of set X. This value can be interpreted as the amount of information learned about set Y by knowing set X and has a minimum value of zero and a maximum value of one. NMI balances the two problematic cases identified above. In the case of each point being assigned its own class, the entropy of the assigned labels is increased, reducing the NMI. In the case of a single label being assigned to the entire set, the mutual information in the denominator is reduced, also reducing the NMI.

While NMI is slightly less understandable than a simple percentage score, it is a well-performing metric for unsupervised machine learning classification. In the case of supervised learning NMI is unnecessary since the resultant classifications are based on, and assigned labels, that exist in the truth or learning set. This allows the use of percentage scores.

4.3.4 Feature sets

In order to explore the limits of shape in this context, and to allow comparison to other calculable properties of these cells, four aspects of the cell were quantified using five sets of features. These are shown in Table 4.1, along with the number of dimensions utilized, the property being quantified, and whether or not the features are local or nonlocal. A local feature is one that is a property of a given cell alone. A nonlocal feature is one which can only be calculated in the context of the given cell's position with respect to other cells and the entire cotyledon as a whole.

The irreducible affine Cartesian moment invariants are calculated using moments up to the 4^{th} order, as described in [107] and defined in [140]. The Zernike moment invariants are

calculated from the central Cartesian moments as described in Equation 3.25 from Chapter III. The Zernike invariants are calculated using moments up to the 20^{th} order and therefore capture more complete level of information about the shape of each cell when compared to the Cartesian invariants.

The volume metric represents the size of each cell and is measured in simple voxels. The volume-normalized second order moments capture orientation information of each cell, but are volume invariant. Nearest neighbors measures the number of cells in close proximity to a given cell. This feature set was calculated for a given cell by dilating the cell through a simple convolution with a spherical kernel of radius 20 voxels and counting the number of cells which come into contact with the dilated region.

Feature set	Abbrev.	Dimensionality	Quantified property	Local/nonlocal
irreducible affine Cartesian moment invariants	C	242	shape	local
Zernike moment invariants	Z	121	shape	local
volume	V	1	size	local
nearest neighbors	N	1	number of nearby cells	nonlocal
volume-normalized second order moments	S	6	orientation	nonlocal

Table 4.1: Properties of feature sets

Every combination of feature sets in Table 4.1, with the exception of those which contain both Zernike and Cartesian invariants, were tested for a total of 23 datasets. In this chapter, datasets will be referred to by a short abbreviation indicating the data present in each. For example, the dataset consisting of the Cartesian moment invariants, the volume, and the number of nearest neighbors for each cell would be referred to as CVN, and have a dimensionality of 244. This data is organized into matrices of size $m \times n$ where m is the number of dimensions or features utilized for each cell and n is the number of examined cells.

In order to place equal importance on varying dimensions of the data, each feature set is normalized prior to analysis. This is important due to the large scale differences present in the data. While volume can be on the scale of 10^3 voxels, while some moment invariants

are on the order of 10^{-5}. In both unsupervised and supervised analysis, this normalization takes place on the entire data set prior to processing. Initially, two normalization methods were utilized, that of equation 4.8, a simple normalization with results lying between 0 and 1, and that of equation 4.9, which is the Z-score of the data. No significant differences were found between the results of analyses utilizing the two methods, so the latter method, that represented by equation 4.9, was utilized since preprocessing of the data to this form can often improve PCA results, thus facilitating unsupervised analysis.

$$N(X) = \frac{X - \min(X)}{\max(X) - \min(X)} \tag{4.8}$$

$$Z(X) = \frac{X - \mu(X)}{\sigma(X)} \tag{4.9}$$

4.3.5 Unsupervised machine learning

Data processing and flow

In order to determine parameters of the t-SNE and HDBSCAN algorithms that result in the best performance, a grid-search method was utilized. The perplexity values for t-SNE were restricted to 10, 30, and 50, which represent the recommended minimum, and maximum values from the author of the algorithm [131] as well as a central value. For HDBSCAN, minimum cluster size values were restricted (somewhat arbitrarily after finding that lower values on the scale of 10^2 tended to perform better) to the set {25,50,75,100,125}, and minimum number of samples to {10,35,60}, once again based on minimums and maximums from the originator of the method [138]. Because t-SNE is not a deterministic algorithm, this was repeated for five iterations. Average and maximum values are reported below.

The success of these methods is scored using the normalized mutual information metric discussed in Section 4.3.3. In some cases a percentage score may also be reported but for reasons also discussed in 4.3.3, the percentage is inherently difficult to calculate, and is only shown as a rough estimate. Any comparisons of success should be made using NMI alone.

4.3.6 Supervised machine learning

Data processing and flow

In this work, linear SVM, as discussed in 4.2.1 was utilized. As such, only a single parameter, the error penalty parameter (C), was required to be optimized. Based upon [141], a grid-search technique was undertaken over the range $C = \{2^{-5}, 2^{-4} \dots 2^{20}\}$. Success was quantified using five-fold cross validation. Using this method, the training data is randomly separated into five subsets, and then, to prevent overtraining, for each subset an SVM model is trained on the other four and then tested on the given subset. The classification of each subset is compiled into an overall set and the percentage correct classification is measured.

4.4 Results

4.4.1 Mapping and clustering

Success of clustering

Table 4.2 lists the results of the mapping and clustering procedure on data set 7x and 8x respectively. For each feature set, the included subsets are listed with an 'X' and those excluded with a '-'. The mean NMI is listed, as calculated over five iterations of the algorithm and with every combination of t-SNE and HDBSCAN parameters. Maximum NMI is the single largest NMI value calculated from the data set with any combination of algorithm parameters. The standard deviation of NMI for each cotyledon is listed in the final two columns.

The top five performing feature sets in terms of mean NMI for dataset 7x are, in descending order, CVNS, CNS, VNS, CS, and CN. These feature sets for 8x are CS, CNS, CVNS, VS, and CVS. In both cases, the Cartesian invariants and the second order moments are the most common, indicating that these features capture aspects of the data that allow better differentiation along the lines of cell identity. In terms of maximum NMI, the top feature sets for 7x are CVNS, CNS, CS, ZS, and CVN; for 8x: CVNS, CNS, VS, NS, and CVS. Once again, the Cartesian features and orientation have have the largest success rate in both cases. In both specimens, it appears that the Cartesian invariants out performed the Zernike invariants, and that, at least in the case of the average NMI, comparable re-

Feature set	Cart	Zern	Vol	NN	Sec	NMI (mean) 7x	NMI (max) 7x	NMI (mean) 8x	NMI (max) 8x	NMI σ 7x	NMI σ 8x
C	X	-	-	-	-	0.141	0.315	0.192	0.276	0.037	0.045
CV	X	-	X	-	-	0.202	0.324	0.167	0.291	0.016	0.087
CN	X	-	-	X	-	0.271	0.372	0.178	0.308	0.069	0.042
CS	X	-	-	-	X	0.283	0.415	0.326	0.449	0.015	0.044
CVN	X	-	X	X	-	0.240	0.374	0.262	0.357	0.024	0.055
CVS	X	-	X	-	X	0.234	0.360	0.296	0.412	0.019	0.063
CNS	X	-	-	X	X	0.303	0.497	0.322	0.461	0.034	0.104
CVNS	X	-	X	X	X	0.353	0.627	0.319	0.508	0.035	0.141
Z	-	X	-	-	-	0.076	0.324	0.135	0.314	0.043	0.085
ZV	-	X	X	-	-	0.073	0.324	0.102	0.309	0.005	0.080
ZN	-	X	-	X	-	0.060	0.335	0.065	0.333	0.033	0.029
ZS	-	X	-	-	X	0.162	0.388	0.080	0.321	0.041	0.030
ZVN	-	X	X	X	-	0.055	0.330	0.052	0.283	0.088	0.024
ZVS	-	X	X	-	X	0.168	0.373	0.061	0.304	0.045	0.033
ZNS	-	X	-	X	X	0.095	0.372	0.094	0.332	0.008	0.139
ZVNS	-	X	X	X	X	0.079	0.341	0.055	0.333	0.073	0.119
V	-	-	X	-	-	0.127	0.203	0.124	0.201	0.011	0.017
N	-	-	-	X	-	0.134	0.174	0.161	0.189	$<$ 0.001	$<$ 0.001
S	-	-	-	-	X	0.139	0.245	0.220	0.359	0.089	0.049
VN	-	-	X	X	-	0.215	0.347	0.181	0.306	0.024	0.052
VS	-	-	X	-	X	0.246	0.340	0.302	0.444	0.061	0.032
NS	-	-	-	X	X	0.237	0.316	0.258	0.439	0.061	0.014
VNS	-	-	X	X	X	0.288	0.350	0.247	0.301	0.003	0.025

Table 4.2: Results of unsupervised clustering on both cotyledon data sets. Standard deviation of NMI calculated over five iterations for each feature set for the set of clustering parameters with highest average NMI.

sults were obtained by the VNS feature set in sample 7x and the VS data set in 8x, which exclude shape information in the form of moment invariants, to those of CVNS

Examining the standard deviation of the NMI, it should be noted that values are relatively large in the case of the best performing 7x feature sets, CVNS, and CNS with values of 0.14 and 0.10 respectively. This is also true for the lower performing feature sets of ZNS and ZVNS. Standard deviations values in the 8x set are somewhat less, with the largest values belonging to the feature sets S, ZVN, ZVNS, and CN, having values of 0.090, 0.088, 0.073, 0.07, respectively. None of these feature sets are among the best performing for the 8x dataset.

The observed standard deviations indicate that there is a large degree of variation in the results of t-SNE mapping and HDBSCAN clustering, especially in the case of specimen 7x. In order to examine the relative influence of algorithm parameters and feature sets on the success rate, analysis of variance (ANOVA) [142] and multiple linear regression were performed..

ANOVA

The results of ANOVA tests are shown in Figures 4.3 and 4.4. The ANOVA table for the 7x data set of 4.3, has large F values and correspondingly low p-values for the main effects of all feature sets and algorithm parameters, with the exception of the volume feature set and the min_clust parameter for HDBSCAN. If an α value of 0.05 is utilized as the cutoff for statistical significance, it can be concluded from this evidence that the null hypothesis: the levels of the given feature have no influence on the result of the clustering, can be rejected in all instances except for those previously mentioned. There are several statistically significant interactions. The above holds true in the case of the 8x cotyledon as well, except that the volume term, the effects of which were insignificant in the 7x cotyledon, has statistical significance in this case. The histograms of residuals in Figures 4.3 and 4.4 are very similar in shape for both data sets. The distribution appears to be bimodal, with a sharp peak in the negative region which cuts off at zero, and a lower peak in the positive region.

Analysis of Variance

Source	Sum Sq.	d.f.	Mean Sq.	F	Prob>F
Cart	1.84	1	1.83962	259.06	0
Zern	8.81	18	0.48943	68.92	0
vol	0.009	1	0.00903	1.27	0.2595
NN	1.01	1	1.01048	142.3	0
sec	4.359	1	4.35908	613.85	0
perp	0.655	2	0.32768	46.14	0
min_clust	0.016	2	0.0079	1.11	0.3288
min_samp	0.321	2	0.16064	22.62	0
Cart*vol	0.001	1	0.00127	0.18	0.6718
Cart*NN	1.566	1	1.56645	220.59	0
Cart*sec	0.188	1	0.1882	26.5	0
Cart*perp	0.065	2	0.03268	4.6	0.01
Cart*min_clust	0.009	2	0.00448	0.63	0.5324
Cart*min_samp	0.042	2	0.02098	2.95	0.0521
Zern*perp	3.438	36	0.09549	13.45	0
Zern*min_clust	0.253	36	0.00702	0.99	0.4891
Zern*min_samp	3.428	36	0.09521	13.41	0
vol*NN	0.001	1	0.0006	0.08	0.7709
vol*sec	0.166	1	0.1656	23.32	0
vol*perp	0.131	2	0.06558	9.23	0.0001
vol*min_clust	0.006	2	0.00294	0.41	0.6608
vol*min_samp	0.022	2	0.01107	1.56	0.2104
NN*sec	0.027	1	0.02663	3.75	0.0528
NN*perp	0.006	2	0.00301	0.42	0.6541
NN*min_clust	0.002	2	0.00107	0.15	0.8599
NN*min_samp	0.173	2	0.08666	12.2	0
sec*perp	0.048	2	0.02406	3.39	0.0338
sec*min_clust	0.027	2	0.01339	1.88	0.1519
sec*min_samp	0.014	2	0.007	0.99	0.3734
perp*min_clust	0.068	4	0.01697	2.39	0.0486
perp*min_samp	3.397	4	0.84922	119.59	0
min_clust*min_samp	0.557	4	0.13926	19.61	0
Error	151.158	21286	0.0071		
Total	242.015	21464			

Constrained (Type III) sums of squares.

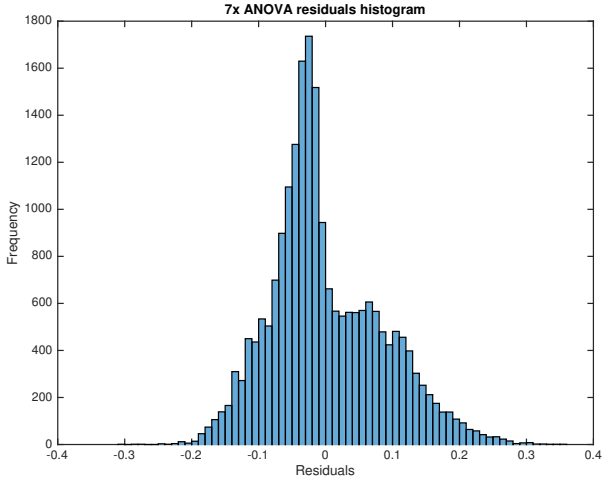

Figure 4.3: ANOVA table and histogram of residuals for the influence of feature sets and algorithm parameters on unsupervised clustering results on 7x data set.

Analysis of Variance

Source	Sum Sq.	d.f.	Mean Sq.	F	Prob>F
Cart	0.865	1	0.86484	149.16	0
Zern	17.504	18	0.97244	167.71	0
vol	0.037	1	0.03712	6.4	0.0114
NN	0.235	1	0.23523	40.57	0
sec	4.958	1	4.95799	855.09	0
perp	0.651	2	0.32568	56.17	0
min_clust	0	2	0.00014	0.02	0.9755
min_samp	0.563	2	0.2813	48.52	0
Cart*vol	0.001	1	0.00136	0.24	0.6278
Cart*NN	0.006	1	0.00578	1	0.3181
Cart*sec	0.682	1	0.68176	117.58	0
Cart*perp	0	2	0.00018	0.03	0.9691
Cart*min_clust	0.002	2	0.00092	0.16	0.8529
Cart*min_samp	0.171	2	0.08551	14.75	0
Zern*perp	8.219	36	0.22829	39.37	0
Zern*min_clust	0.075	36	0.00209	0.36	0.9998
Zern*min_samp	1.619	36	0.04498	7.76	0
vol*NN	0.043	1	0.04295	7.41	0.0065
vol*sec	0.021	1	0.02059	3.55	0.0595
vol*perp	0.078	2	0.03899	6.73	0.0012
vol*min_clust	0.007	2	0.00344	0.59	0.5528
vol*min_samp	0.058	2	0.02922	5.04	0.0065
NN*sec	0.115	1	0.11521	19.87	0
NN*perp	0.059	2	0.02928	5.05	0.0064
NN*min_clust	0	2	0.00019	0.03	0.9679
NN*min_samp	0.069	2	0.03445	5.94	0.0026
sec*perp	0.012	2	0.00604	1.04	0.353
sec*min_clust	0	2	0.00005	0.01	0.9917
sec*min_samp	0.046	2	0.02285	3.94	0.0195
perp*min_clust	0.025	4	0.00619	1.07	0.3704
perp*min_samp	1.018	4	0.25457	43.9	0
min_clust*min_samp	0.097	4	0.02419	4.17	0.0022
Error	123.421	21286	0.0058		
Total	202.151	21464			

Constrained (Type III) sums of squares.

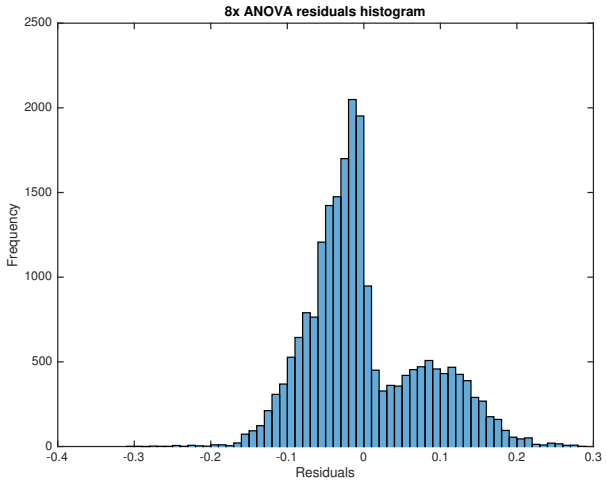

Figure 4.4: ANOVA table and histogram of residuals for the influence of feature sets and algorithm parameters on unsupervised clustering results on 8x data set.

Multiple linear regression

Table 4.3 shows the results of multiple linear regression conducted on the full set of features and algorithm parameters to the unsupervised analysis. The coefficient of multiple regression is 0.459 and 0.350 for datasets 7x and 8x respectively, indicating that the variance in the independent variables accounts for about this proportion of variance in the result [143]. The F-value for the linear models is 454 and 288 respectively, far above the F-value of significance: approximately 3.5 for both specimens [144].

In data set 7x, the correlation coefficients for Cartesian invariants and the second order moments is positive, while the other values are negative. In the 8x data set, the correlation coefficient for the second order moments is negative, indicating that the quality of the clustering results in the different cotyledons may be influenced differently by the same feature sets. In order to measure the relative importance of these feature sets, a ΔR_m^2 value is given in each table for each parameter. This value is the squared semipartial correlation and is the amount by which R_m^2 decreases when these features are removed from the multiple linear correlation analysis [145]. This value is largest for the second order moments and Cartesian invariants and smallest for volume for the 7x cotyledon. For the 8x cotyledon the largest squared semipartial correlation is that of the Cartesian invariants while the second order moments has the smallest value.

The analysis summarized in Table 4.3 also includes the influence of t-SNE and HDBSCAN parameters. The linear correlation coefficients for individual parameters indicate a negative correlation for all three values: perplexity, minimum number of clusters, and minimum number of samples, suggesting that the lower these values are, the better the algorithm performs. The semipartial correlation coefficient for these parameters is the largest in the case of perplexity, which indicates that out of these values, perplexity has the greatest influence over the final results, although compared overall to the feature values, this influence is small.

MLR Dataset 7x	Cartesian	Zernike	Volume	NN	secord	perp	min_clust	min_samp		const	R_m^2
coefficient	0.0789	-0.0060	-0.0130	0.0078	0.1163	-0.0009	>-0.0001	-0.0003		0.139	0.276
R_l	0.2978	-0.0552	-0.0260	-0.1082	0.3489	-0.1394	-0.0332	-0.0586			
ΔR_m^2	0.0164	0.0302	0.0030	0.0008	0.1297	0.0194	0.0011	0.0034			
F-value	1025.69										

MLR Dataset 8x	Cartesian	Zernike	Volume	NN	secord	perp	min_clust	min_samp		const	R_m^2
coefficient	0.0469	-0.0097	0.0061	0.0300	0.0642	-0.0010	> -0.0001	-0.0002		0.1743	0.2676
R_l	0.3457	-0.3778	-0.1192	-0.0967	-0.0203	-0.1644	-0.0050	-0.0517			
ΔR_m^2	0.0069	0.0944	0.0008	0.0143	0.0473	0.0270	<0.0001	0.0027			
F-value	979.91										

Table 4.3: Results of multiple linear correlation analysis performed on variables affecting clustering results. Dependent variable is the mean mutual information for five iterations. r_l is the linear correlation coefficient for the given parameter, while R_m^2 is the multiple linear correlation coefficient. ΔR_m^2 is the squared semipartial correlation for the given feature.

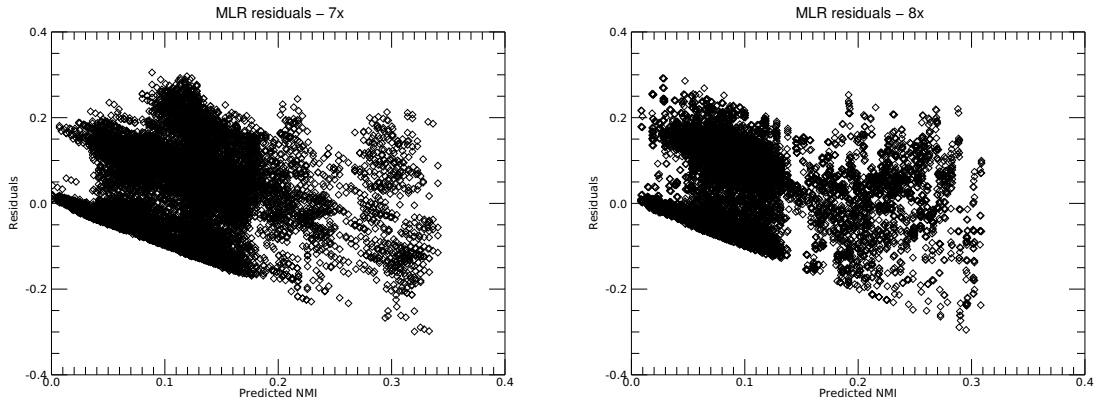

Figure 4.5: Residual plots for 7x and 8x respectively after interaction-free multiple linear regression. There is a clear bimodal distribution of points and a negative trend, indicating that some factor not listed is likely at play, or interactions, which have not been included, are having an effect.

Cluster identification

Figure 4.6 illustrates a typical set of mapped data after application of t-SNE and HDBSCAN. In this case, the feature sets used were Cartesian invariants, volume, nearest neighbors, and the second order moments. Perplexity was set at 30, with a minimum cluster size of 50 and a minimum sample size of 35. In this specific example, the number of cell clusters was larger than the number of cell classes. In order to account for this discrepancy, clusters were merged by combining two clusters which shared the two closest points of data. This was repeated until the desired number of clusters was achieved. This merging mimics the behavior of HDBSCAN and performs better than merging by nearest centers of mass for each cluster. This process resulted in the second set of mapped data in Figure 4.6.

3D representations of cotyledons, colored by assigned class, are shown in Figure 4.8. Because clusters are assigned arbitrary labels by the HDBSCAN, it is difficult to quantitatively compare identified clusters with specific classes of cells. However, it is clear that in this instance the algorithm has made a distinction between epidermal cells (red), vascular cells and surrounding spongy mesophyll (purple), spongy mesophyll (yellow) and palisade mesophyll (green). However, it is evident that the number of mislabeled cells is significant, which corresponds to the NMI for this clustering of 0.42.

Figure 4.7 displays the results of a different instance of t-SNE mapping and HDBSCAN clustering undertaken with the same feature sets as Figure 4.6, but with a perplexity of 10, minimum cluster size of 50 and a minimum sample size of 35. This generated four total clusters with an NMI of 0.346. Despite the use of identical features, the identified clusters differ to a large degree, as shown in Figure 4.9. The purple and red clusters comprise the vascular system of the cell, with the primary differentiating feature possibly being orientation. The yellow cluster is primary composed of the top layer of the epidermis. The final green cluster contains all of the mesophyll, and the bottom layer of the epidermis.

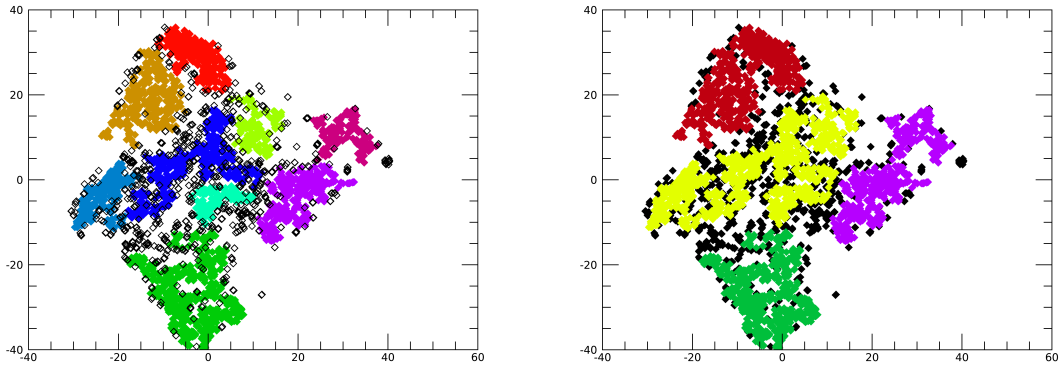

Figure 4.6: Similar cells identified after initial t-SNE and HDBSCAN clustering (left) and tree merging (right) of 7x dataset with CVNS feature set, perplexity of 30, min_clust of 50, and min_samp of 35. Clustering shown has an NMI 0f 0.420.

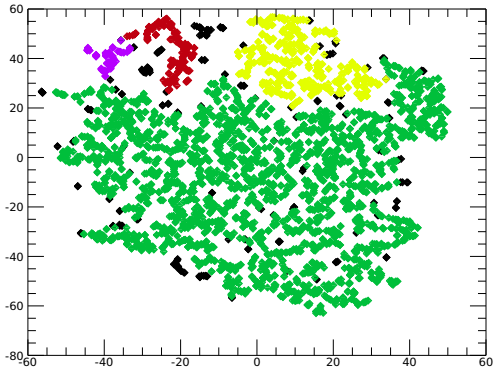

Figure 4.7: Generated t-SNE mapping and HDBSCAN clustering on 7x dataset with CVNS features, perplexity of 10, min_clusts of 50, and min_samples of 10. Clustering shown has an NMI 0f 0.346.

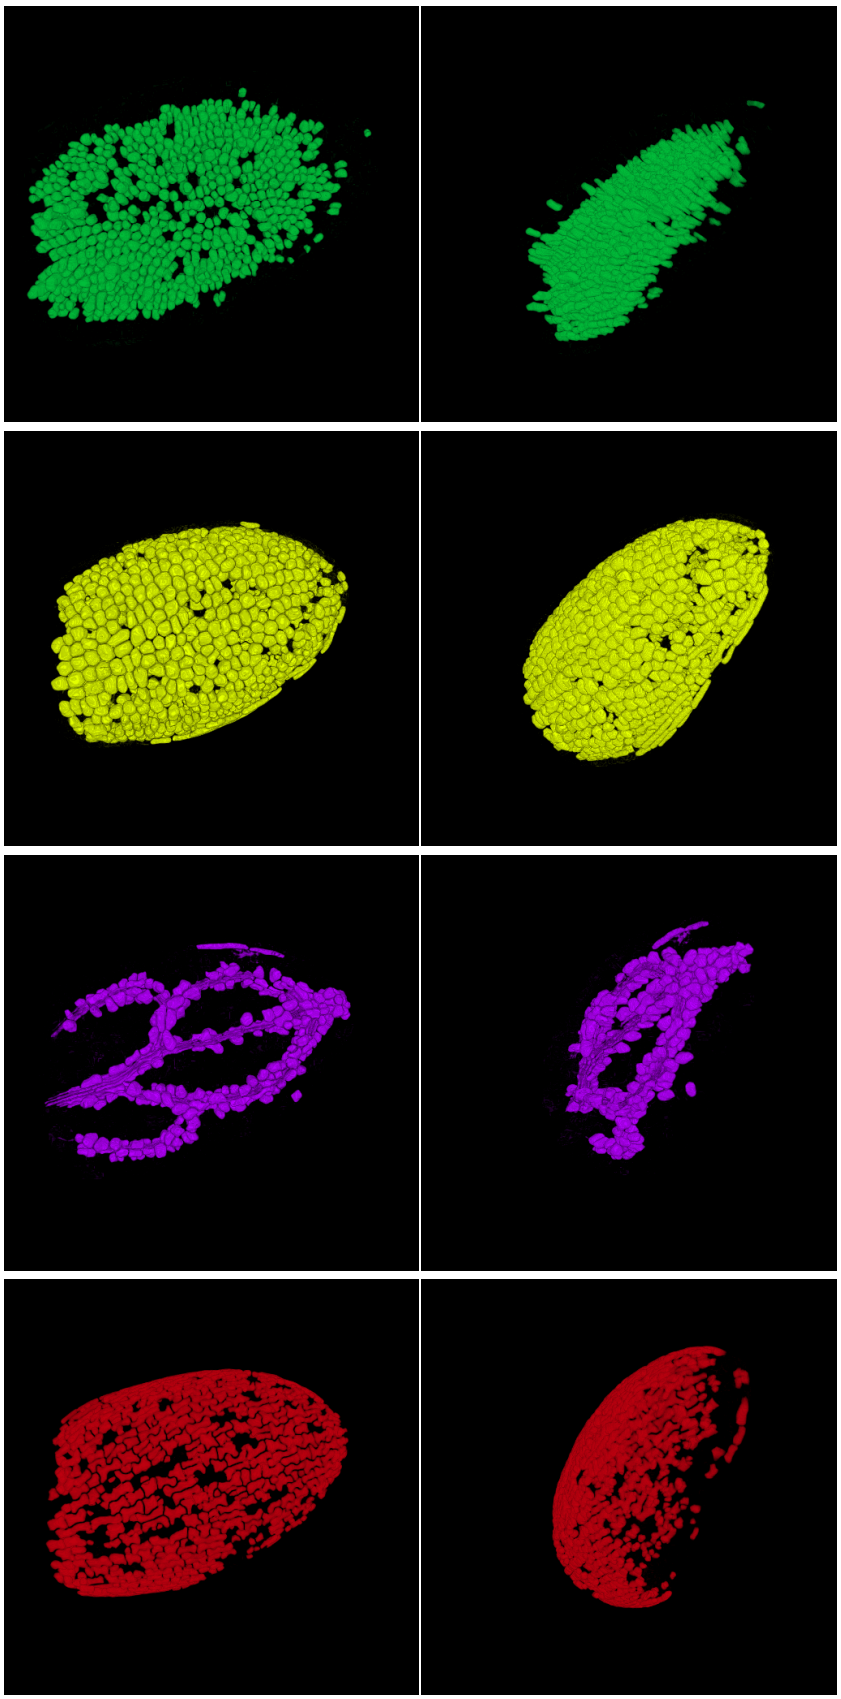

Figure 4.8: Identified clusters based upon t-SNE mapping and HDBSCAN clustering of Figure 4.6.

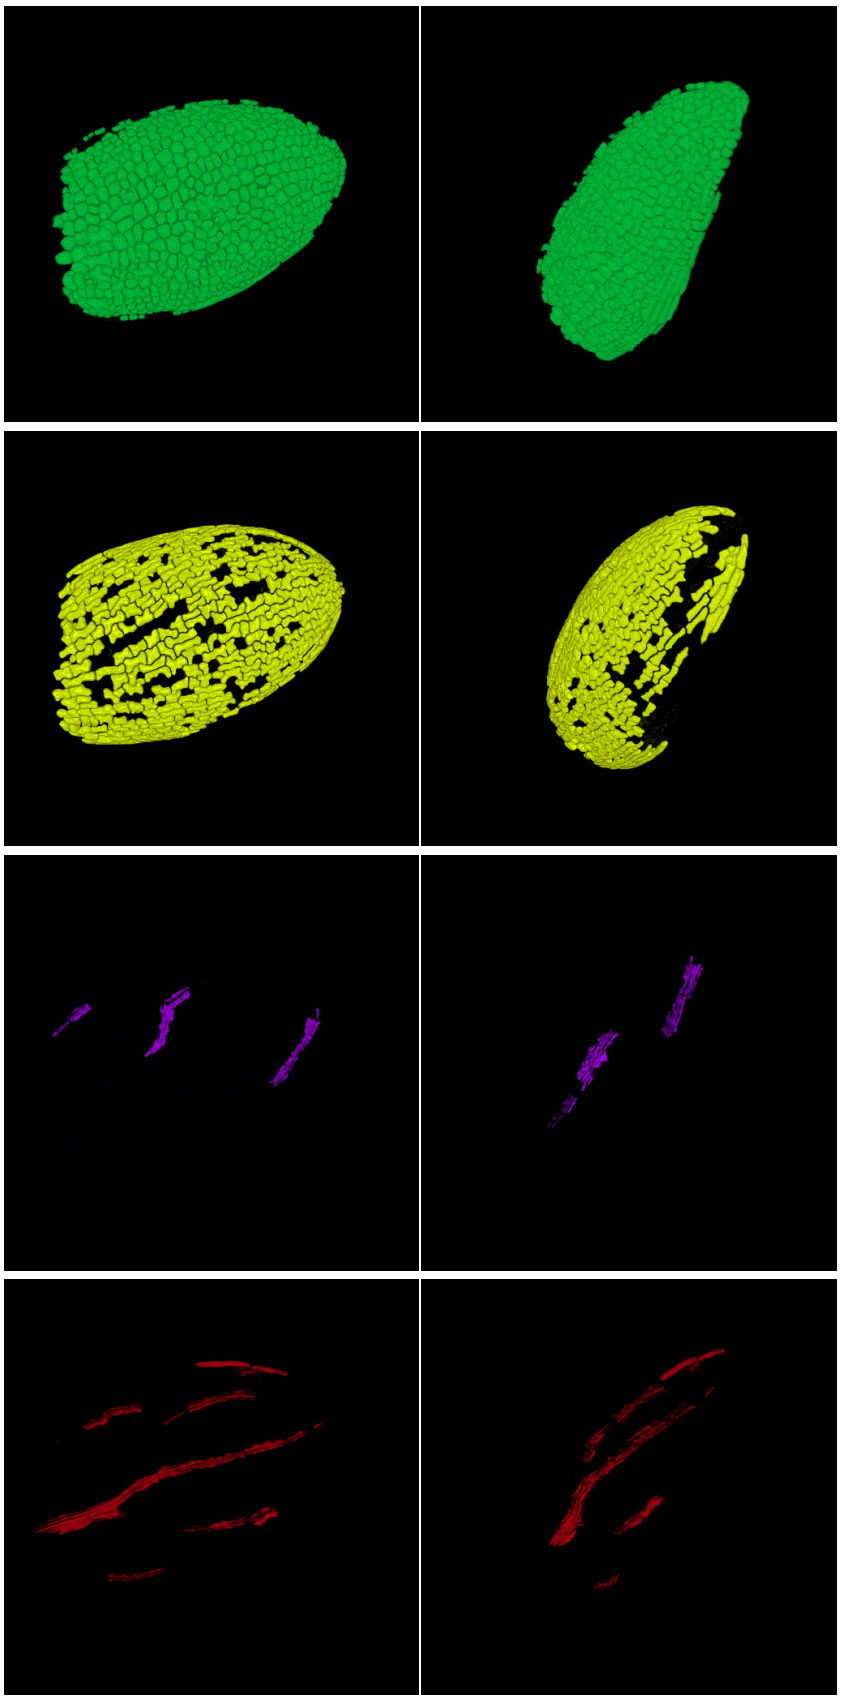

Figure 4.9: Identified clusters based upon t-SNE mapping and HDBSCAN clustering of Figure 4.7.

4.4.2 Support vector machine analysis

A summary of support vector machine analysis described in Section 4.2.1 is shown in Table 4.4. The feature sets with the largest values of percentage correct classification for the 7x data set are VNS, CVNS, ZVNS, CNS, and ZNS, while the worse performing feature sets were V, N, and C. Similar feature sets scored well in the 8x data set, with the highest values belonging to CVNS, ZVNS, VNS, ZNS, and CNS, while the feature sets which performed the worse were V, N, and C, as with the 7x set.

Three line plots illustrating the full results of the grid-search over values of the error penalty parameter (C) are shown in Figure 4.10. The general shape of these plots, regardless whether they include Cartesian or Zernike invariants or exclude them all together, are similar. In addition, the slopes of the lines from $C = 2^{-5}$ to about $C = 2^{12}$ are roughly horizontal, indicating that there is little variation in the results with changing error parameter, although the poorer-performing feature sets showed some improvement with increasing C. Greater variation in the success rate in the region beyond this parameter indicates that the results are variable and possibly dependent on random initialization of the SVM algorithm.

Feature set	Cartesian	Zernike	Volume	NN	2nd ord	7x %	7x NMI	8x %	8x NMI
C	X	-	-	-	-	65.31	0.373	60.64	0.388
CV	X	-	X	-	-	66.97	0.362	63.05	0.370
CN	X	-	-	X	-	76.21	0.458	75.81	0.452
CS	X	-	-	-	X	81.53	0.592	80.50	0.585
CVN	X	-	X	X	-	77.15	0.481	77.99	0.503
CVS	X	-	X	-	X	84.75	0.627	84.60	0.608
CNS	X	-	-	X	X	89.62	0.699	91.29	0.723
CVNS	X	-	X	X	X	90.60	0.720	92.55	0.758
Z	-	X	-	-	-	71.06	0.393	67.96	0.388
ZV	-	X	X	-	-	70.62	0.394	68.42	0.396
ZN	-	X	-	X	-	80.91	0.519	78.82	0.494
ZS	-	X	-	-	X	83.12	0.601	84.32	0.617
ZVN	-	X	X	X	-	81.22	0.525	78.75	0.501
ZVS	-	X	X	-	X	83.37	0.604	84.65	0.627
ZNS	-	X	-	X	X	90.49	0.715	92.18	0.736
ZVNS	-	X	X	X	X	90.30	0.711	92.12	0.751
V	-	-	X	-	-	44.86	0.196	46.89	0.219
N	-	-	-	X	-	51.35	0.218	50.71	0.221
S	-	-	-	-	X	80.56	0.564	81.06	0.580
VN	-	-	X	X	-	56.74	0.288	57.71	0.312
VS	-	-	X	-	X	84.08	0.609	85.52	0.627
NS	-	-	-	X	X	89.16	0.684	91.57	0.711
VNS	-	-	X	X	X	90.52	0.717	92.68	0.752

Table 4.4: Summary of support vector machine analysis of 7x and 8x datasets via five-fold cross validation.

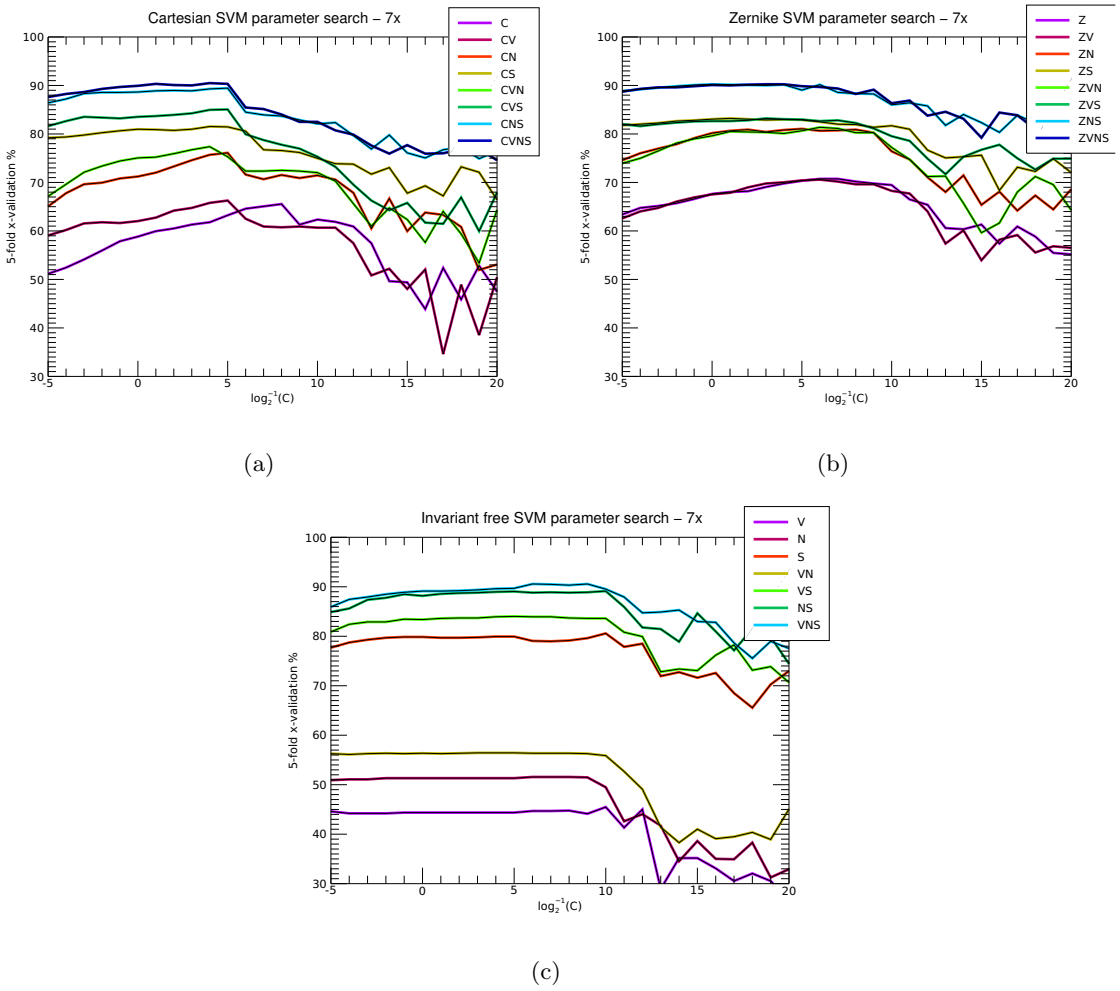

Figure 4.10: Results of SVM grid search as measured by five-fold cross validation percentage for 7x data set.

Analysis of Variance

Source	Sum Sq.	d.f.	Mean Sq.	F	Prob>F
Cart	1653	1	1652.96	16854.41	0
Zern	2367.4	1	2367.36	24138.79	0
vol	14.1	1	14.09	143.66	0
NN	1521.1	1	1521.12	15510.18	0
sec	3401.6	1	3401.65	34684.98	0
Cart*vol	16.7	1	16.65	169.78	0
Cart*NN	3.7	1	3.73	38.07	0
Cart*sec	1505.9	1	1505.87	15354.58	0
Zern*vol	48.8	1	48.8	497.57	0
Zern*NN	4.2	1	4.22	43.01	0
Zern*sec	2130.4	1	2130.36	21722.26	0
vol*NN	2.2	1	2.2	22.44	0
vol*sec	1	1	1	10.22	0.0019
NN*sec	55.3	1	55.29	563.76	0
Cart*vol*NN	0	1	0	0	1
Cart*vol*sec	12.2	1	12.25	124.87	0
Cart*NN*sec	1	1	1	10.19	0.0019
Zern*vol*NN	1.9	1	1.9	19.37	0
Zern*vol*sec	8.5	1	8.52	86.85	0
Zern*NN*sec	2	1	2.01	20.52	0
vol*NN*sec	2.3	1	2.31	23.54	0
Error	9.1	93	0.1		
Total	18506.8	114			

Constrained (Type III) sums of squares.

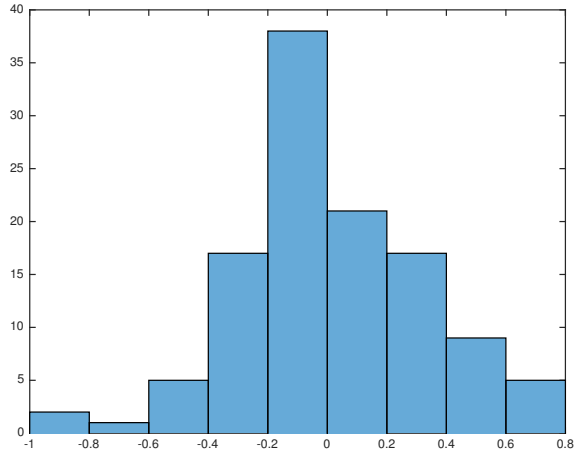

Figure 4.11: ANOVA table and histogram of residuals for the influence of feature sets supervised, SVM classifier on 7x data set.

Analysis of Variance

Source	Sum Sq.	d.f.	Mean Sq.	F	Prob>F
Cart	1215.9	1	1215.93	10148.04	0
Zern	1906.2	1	1906.16	15908.58	0
vol	35.3	1	35.32	294.79	0
NN	2276.7	1	2276.75	19001.51	0
sec	5343.1	1	5343.13	44593.21	0
Cart*vol	24.6	1	24.56	204.94	0
Cart*NN	9.8	1	9.76	81.48	0
Cart*sec	1329.7	1	1329.67	11097.25	0
Zern*vol	78.9	1	78.86	658.19	0
Zern*NN	4.3	1	4.34	36.25	0
Zern*sec	1729.9	1	1729.92	14437.7	0
vol*NN	5	1	5.03	42	0
vol*sec	0.1	1	0.13	1.12	0.2925
NN*sec	92.6	1	92.65	773.21	0
Cart*vol*NN	1	1	1.02	8.51	0.0044
Cart*vol*sec	23	1	22.98	191.75	0
Cart*NN*sec	5	1	4.99	41.61	0
Zern*vol*NN	3.7	1	3.69	30.83	0
Zern*vol*sec	19.4	1	19.43	162.18	0
Zern*NN*sec	0	1	0.01	0.08	0.7778
vol*NN*sec	1.9	1	1.92	16.02	0.0001
Error	11.1	93	0.12		
Total	21142.3	114			

Constrained (Type III) sums of squares.

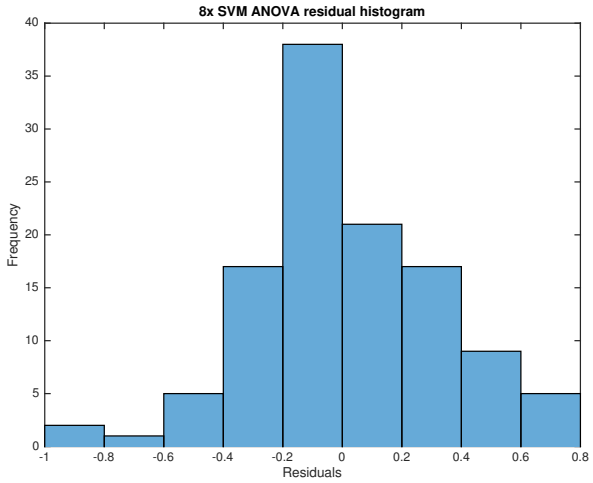

Figure 4.12: ANOVA table and histogram of residuals for the influence of feature sets supervised, SVM classifier on 7x data set.

MLR 7x	Cartesian	Zernike	Volume	NN	secord		const	R_m^2
coefficient	10.01	12.38	0.59	7.77	20.08		54.79	0.806
R_l	0.092	0.227	-0.053	0.242	0.749			
ΔR_m^2	0.100	0.153	0.001	0.093	0.619			
F-value	90.52							

MLR 8x	Cartesian	Zernike	Volume	NN	secord		const	R_m^2
coefficient	8.39	10.99	1.41	9.28	22.76		53.19	0.852
R_l	0.028	0.173	0.002	0.286	0.827			
ΔR_m^2	0.061	0.105	0.003	0.116	0.696			
F-value	125.48							

Table 4.5: Results of multiple linear correlation analysis performed on indicated feature sets for SVM. Dependent variable is the percent success from five-fold cross validation over five iterations. "secord" indicates second order moments (orientation). r_l is the linear correlation coefficient for the given parameter, while R_m^2 is the multiple linear correlation coefficient. ΔR_m^2 is the squared semipartial correlation for the given feature.

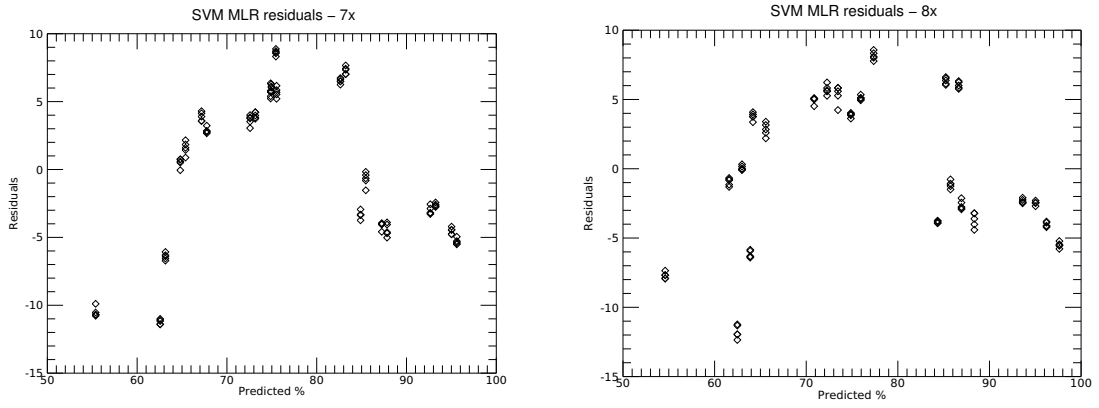

Figure 4.13: Residual plots for the multiple linear regression analysis of Table 4.5

Figure 4.17 shows the result of application of the SVM model for the CVNS feature set from the 7x cotyledon applied to the same cotyledon. Figures 4.11 and 4.12 show ANOVA tables with and without interaction terms. Table 4.5 shows the results of multiple linear correlation on the results of five-fold cross validation SVM. In this model the dependent variable was the maximum percentage found over all searched C values, (*i.e.* the maximum value from a line from Figure 4.10). Figures 4.15 and 4.16 show confusion matrices from various SVM classifiers generated by the indicated feature sets in the 7x and 8x data sets respectively.

Table 4.6 breaks down the generalizability of a model developed on one cotyledon, as

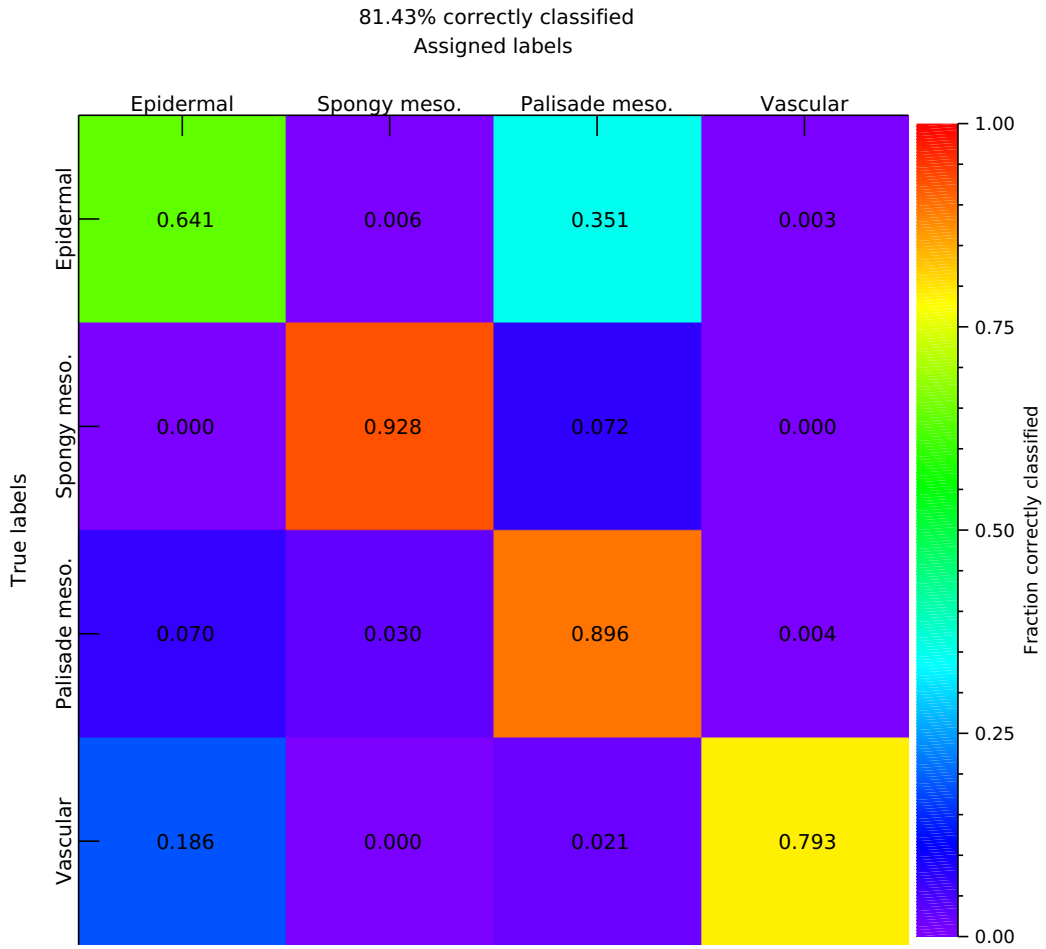

Figure 4.14: Example of confusion matrix, shown here larger to illustrate the position of the categories and axis titles.

tested on the other. It is clear that simple application of a model generated from one cotyledon to another is nearly as successful as the application of one to itself. The final column of the table lists the five-fold cross validation success percentage when a model trained on both cotyledons is applied to both.

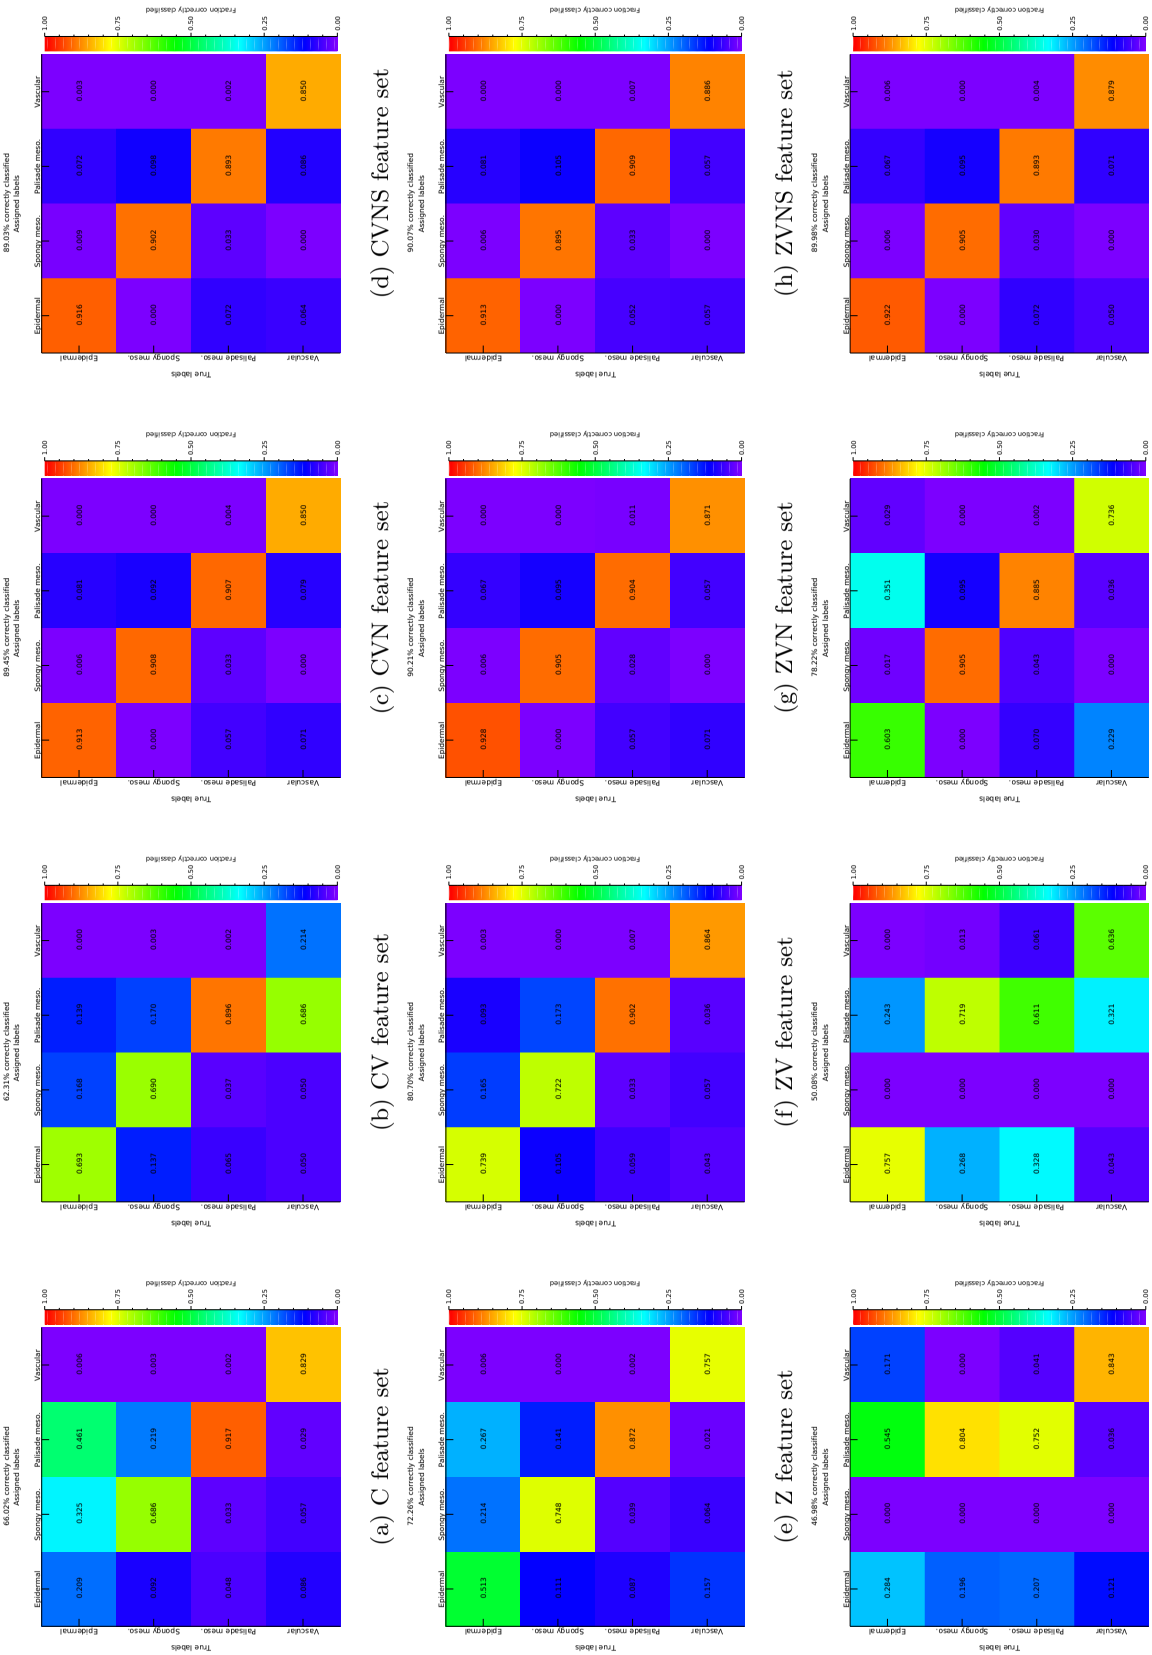

Figure 4.15: Set of confusion matrices for SVM of 7x data set

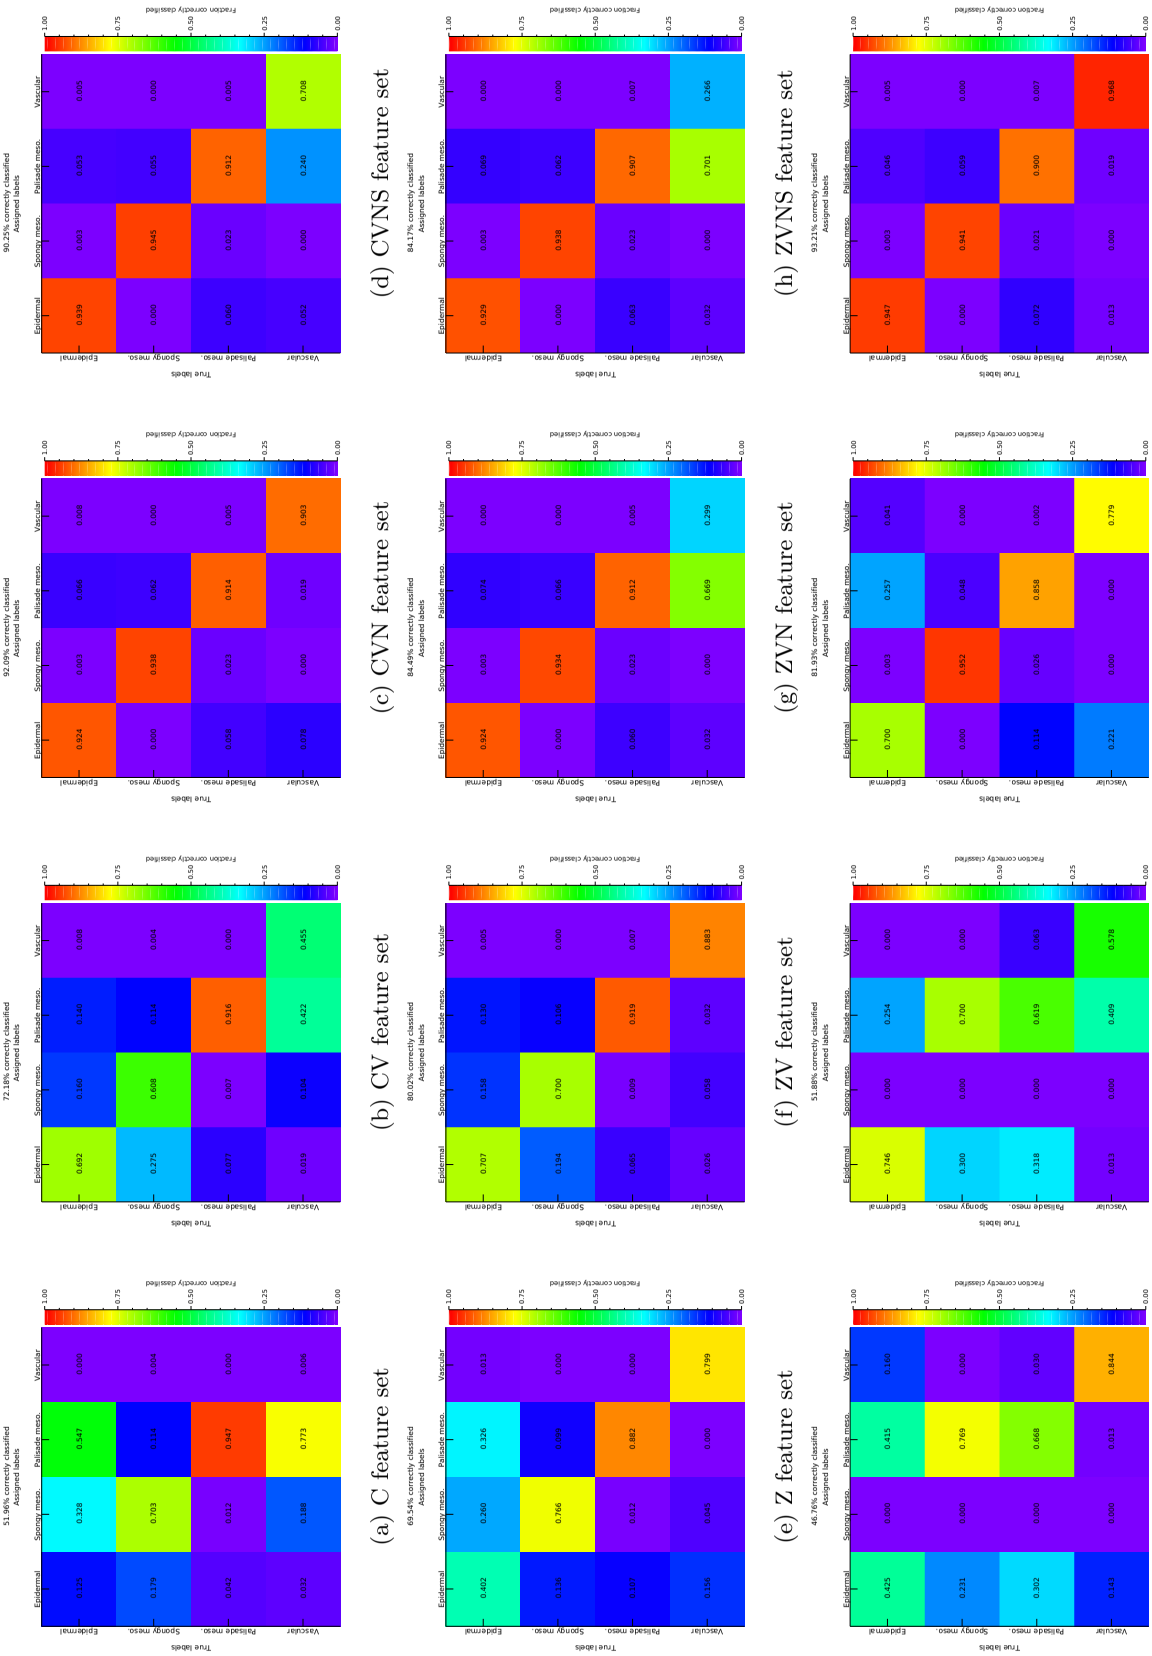

Figure 4.16: Set of confusion matrices for SVM of 8x data set

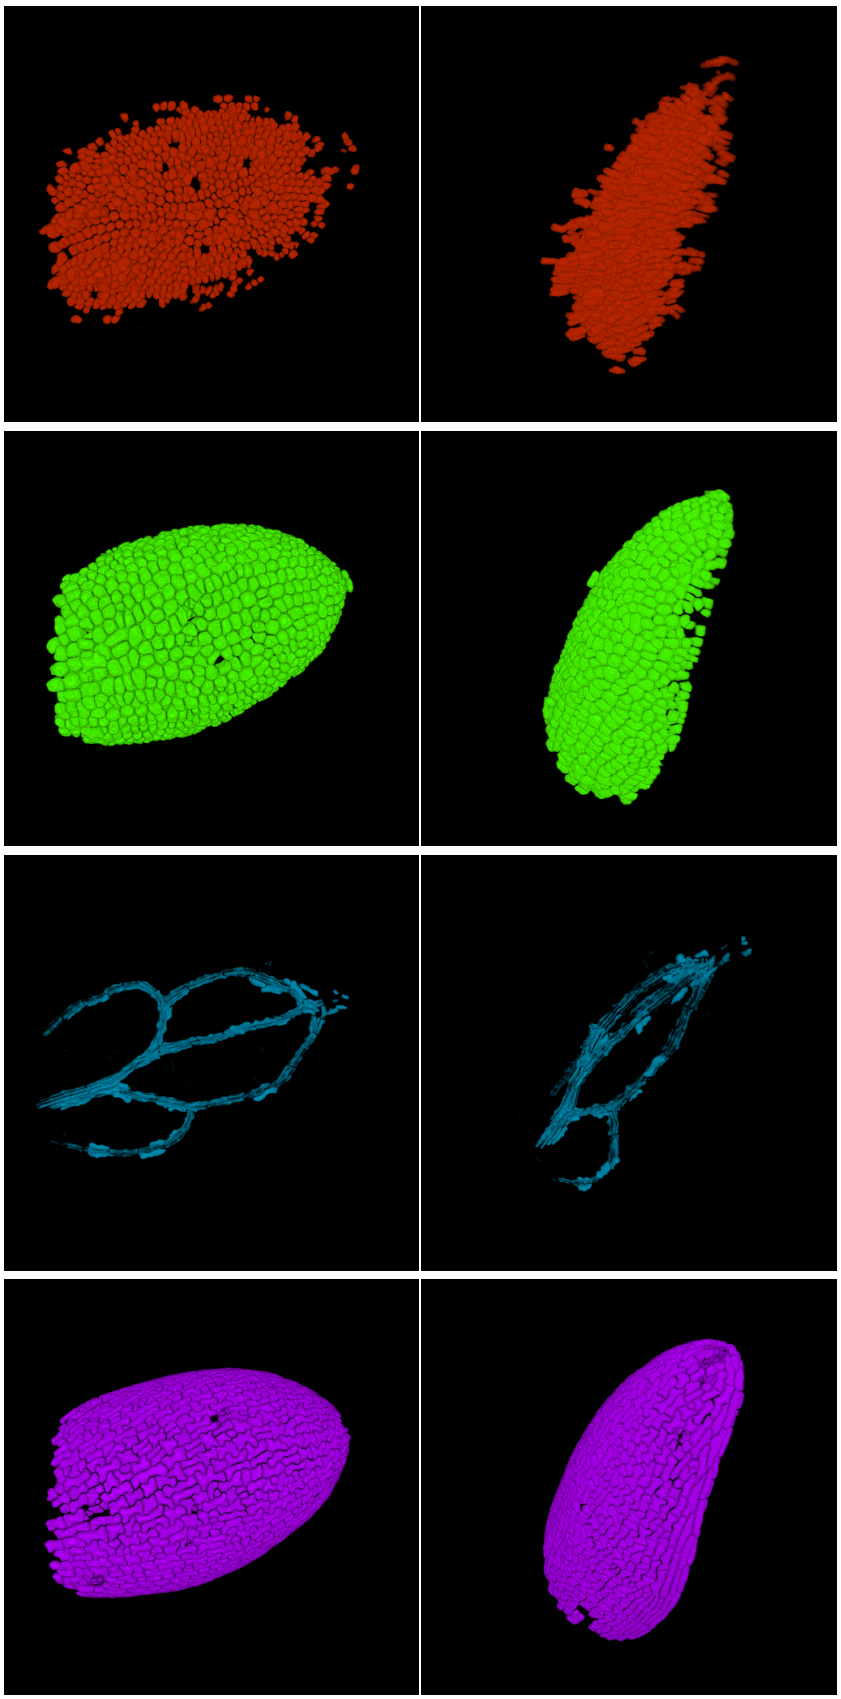

Figure 4.17: Identified clusters based upon SVM classifier for 7x dataset generated by CVNS feature set.

Feature set	Cartesian	Zernike	Volume	NN	2nd ord	7x to 8x %	7x to 8x NMI	8x to 7x %	8x to 7x NMI
C	X	-	-	-	-	51.24	0.368	60.56	0.343
CV	X	-	X	-	-	53.96	0.289	64.68	0.351
CN	X	-	-	X	-	69.38	0.402	75.65	0.440
CS	X	-	-	-	X	73.7	0.570	78.63	0.542
CVN	X	-	X	X	-	68.35	0.398	76.37	0.471
CVS	X	-	X	-	X	74.18	0.522	82.74	0.578
CNS	X	-	-	X	X	88.33	0.672	87.98	0.659
CVNS	X	-	X	X	X	85.45	0.625	89.52	0.695
Z	-	X	-	-	-	67.87	0.409	71.69	0.394
ZV	-	X	X	-	-	68.19	0.419	72.18	0.408
ZN	-	X	-	X	-	78.82	0.500	80.56	0.513
ZS	-	X	-	-	X	76.74	0.551	83.47	0.591
ZVN	-	X	X	X	-	78.58	0.496	80.40	0.513
ZVS	-	X	X	-	X	77.38	0.552	83.55	0.594
ZNS	-	X	-	X	X	90.17	0.709	89.44	0.692
ZVNS	-	X	X	X	X	89.77	0.697	89.76	0.700
V	-	-	X	-	-	48.76	0.236	47.5	0.234
N	-	-	-	X	-	45.24	0.2590	44.35	0.265
S	-	-	-	-	X	81.14	0.5879	80.81	0.564
VN	-	-	X	X	-	57.39	0.323	58.06	0.335
VS	-	-	X	-	X	83.69	0.612	84.03	0.620
NS	-	-	-	X	X	89.77	0.704	89.76	0.700
VNS	-	-	X	X	X	90.49	0.718	89.92	0.722

Table 4.6: Summary of cross-cotyledon application of models

4.5 Discussion

As expected, the results of the two clustering schemes, supervised and unsupervised are quite different. An unsupervised approach allows natural groupings present in the data to emerge and be identified, while a supervised method utilizes the data points to create a classifier that can enforce user-defined groupings upon new points of data. What follows is a discussion of what can be learned from the results of both approaches with regards to the degree of information about a cell's class that can be gained from its shape.

4.5.1 Unsupervised analysis

The tables of features and scores for the 7x and 8x cotyledons (Table 4.2) give an at-a-glance view of the relative strengths of each feature set at achieving a clustering that is close to replicating the groupings of cell type. The mean results are summarized in Figure 4.18. Ultimately, the method was not capable of consistently generating a set of clusterings that closely followed the delineations of epidermal, spongy and palisade mesophyll, and vascular cells against which they were scored: no feature set had an average NMI value of greater than 0.4. However, maximum values of 0.627 and 0.508 were achieved for the 7x and 8x cotyledons respectively. These maximum values match those of some of the trained models, indicating that, at least on occasion, the mapping and clustering method does find natural clusters in the data which correspond to cell identity.

In general, the 7x and 8x cotyledons responded similarly to the feature sets used: feature sets which included Cartesian invariants outperformed similar sets with Zernike invariants or without invariant features. In fact, by comparing feature sets with and without Zernike invariants, such as V and ZV, the addition of these invariants appears to actually *decrease* the quality of the classification. This will be discussed in more detail below, but it may be due to the way in which dimensions are treated during t-SNE and the orthogonal nature of the Zernike invariants.

Statistical analysis

ANOVA tables for the unsupervised analysis are shown in Figures 4.3 and 4.4. The ANOVA tables illustrate the significance of the listed feature sets and their interactions. P-values

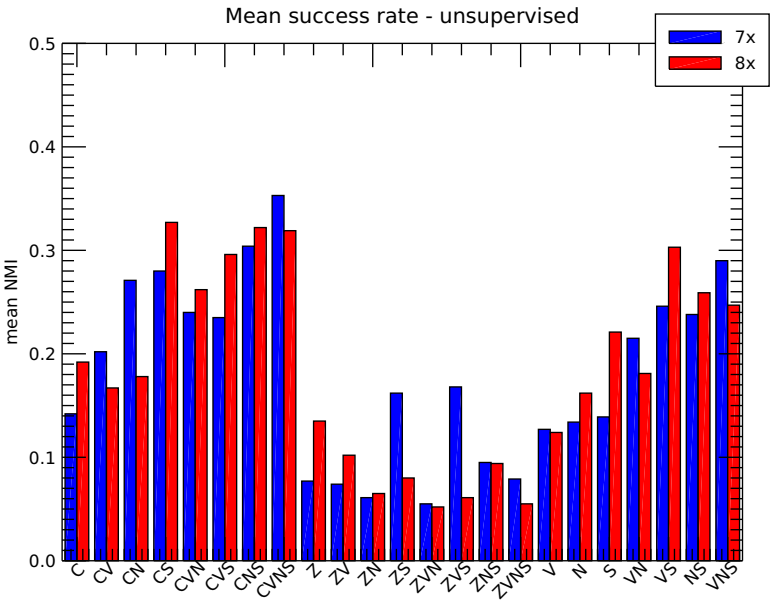

Figure 4.18: Comparison of cross validation success between cotyledons in t-SNE mapping and HDBSCAN clustering with indicated feature sets.Values taken from Table 4.2

less than the α value of 0.05 indicate that *at least one* combination of the listed features generate mean success rates that are significantly different, statistically, than not including them at all. Put another way, for an interaction A*B*C, if the p-value is larger than 0.05, there is no evidence that the combination of features influences the results in a way that cannot be explained by an interaction of A*B, B*C, or A, B, and C alone.

The accompanying residual histograms do not follow a normal distribution. The bimodal distributions indicate that there may be some underlying relationship in the data that has yet to be captured. The large peak just to the negative side of zero in both Figures 4.3 and 4.4 indicates that the model generated by ANOVA is predicting values that are consistently slightly larger than the actual NMI results. It is unclear at this time why this is the case.

Overall, the ANOVA tables indicate that the inclusion of the volume feature in the analysis of the 7x cotyledon had no statistically significant effect on the success of the analysis. In addition, the minimum cluster size, which in this search was limited to $\{25, 75, 125\}$, had no effect on the overall success either. In future analyses, it may be expedient to exclude volume from an analysis. It is also likely that minimum cluster size had no effect in the

given range because even the maximum value is likely less than the number of cells in the smallest sized class. Again, in the 7x cotyledon, many of the interaction terms which lack significance are those interactions with main effects that also lack significance.

Comparing ANOVA tables between the 7x and 8x cotyledon can give an idea of the way in which the cell data is distributed in data space. The volume feature is significant in the 8x cotyledon, in contrast to the 7x data set. This indicates that, at least to some degree, the difference in volume between the classes of cells in the 8x cotyledon may be larger than in the 7x cotyledon. This would lead to this feature contributing more in the t-SNE mapping and subsequent clustering.

The multiple linear regression tables of Table 4.5 and the accompanying residual plots of Figure 4.13 show the results of MLR applied to the mean NMI values resulting from the unsupervised analysis. The coefficients of multiple regression R_m^2 for both 7x and 8x cotyledons indicate that the linear models generated here can only account for approximately 27% of the variation in the data, and thus may be missing a factor which influences the data, similar to the ANOVA model. This is further supported by the residual plots which display some negative correlation, when, if all variables influencing the data would have been accounted for, they would show none.

Nevertheless, these results are somewhat expected. The final NMI values are a scalar quantity which are used to quantify the result of a dimensionality reduction method which is randomly initiated [131] and can result in visually distinct mappings of identical data. Additionally, as previously stated, an unsupervised machine learning method can only identify natural clusters within the data and is not constructed in such a way to deliver a result with any sort of final goal in mind. Similar to the ANOVA tables, there is still some important information that can be gleaned from the MLR analysis.

The F-values are far above the critical F values of 2.01 [144], indicating that the linear model used here is at least better than the mean of the response [146]. The squared semipartial correlation values can be used as a measurement of variable importance, as they quantify the fraction of the variance in the response variable due to the given independent variable [147]. Listing the feature sets in ascending order of importance for the 7x cotyledon gives: nearest neighbors, volume, Cartesian invariants, Zernike invariants, and the second

order moments. The same list for the 8x cotyledon is: volume, Cartesian invariants, nearest neighbors, second order moments and the Zernike invariants.

The MLR tables display some interesting results. For example, in some cases, the nearest neighbors in the 7x table or the second order moments in the 8x table have positive coefficient values, but negative correlation values. The coefficient values indicate the increase in the response variable (mean NMI) with the inclusion of that feature set, all else being constant. This suggests that the inclusion of these features improves the analysis. However, the correlation values are negative, indicating that the inclusion of this feature correlates negatively with the overall results. This is known as Simpson's Paradox [148], and can occur for a couple of reasons. First, the coefficient is calculated with all other variables being held constant, while the correlation is not. If a variable is correlated strongly with another, it could lead to this effect when the variable is excluded for calculation of the coefficient. Alternatively, if there is an underlying variable that is not considered but is also correlated with the variable in question, this effect could be seen as well. It is difficult to say for sure, but there is likely some correlation among the independent variables here. The trend in the residuals also indicates that some other effect may be at play.

However, in the end, the negative coefficients, calculated by holding all other variables constant, suggests that the success of the method is reduced in both data sets with the introduction of the Zernike invariants, which was remarked upon above and can be observed in Figure 4.18. A positive coefficient for the second order moments in the 8x cotyledon is supported by Figure 4.18 and the differences between ZN and ZNS, V and VS, and C and CS. The effects of volume and nearest neighbors are less clear. The inclusion of volume decreases average NMI from C to CV in 7x but decreases NMI in the other cotyledon. Similar effects can be seen with nearest neighbors. Ultimately, the single feature with the greatest positive effect for the 7x data set was the second order moments followed by the Cartesian invariants. In the 8x cotyledon, this was true as well.

The two constructed 3D volumes shown in Figures 4.8 and 4.9 were generated from the same data set, with different perplexity values. It is clear from the differences in the clustering what kind of effect a single change of parameter can elicit. While the first set of clusters has a higher score which more closely matches the cell functions, there are some

errors. The red cluster has only the upper layer of the epidermis, while the lower portion is contained in the yellow cluster. Most obvious is the fact that the purple cluster, which contains the vascular cells, also contains all of the nearby spongy mesophyll cells as well. These mesophyll cells must have slightly distinct morphologies to allow the passage of the vascular cells and these subtle differences were highlighted in this clustering.

The second clustering of Figure 4.9, has clearly isolated the vascular system, but in two separate clusters: the red and purple. It is likely that the second order moments, which contain orientation information were used to make the distinction between these two clusters, as the red cells are all oriented along the long axis of the cotyledon, and the purple are perpendicular. Once again, the epidermal cells of the yellow cluster are strictly limited to the upper region, while the lower region is contained in the same cluster as all of the mesophyll.

In total, the unsupervised analysis shows that there are clear delineations in the data which correspond to the cell identities and therefore the cell functions. Ultimately, the second order moment shape features, which contain shape and orientation information are likely the single most important feature in this mapping and clustering. However, the Cartesian invariants play an important role as well. In both cotyledons the Cartesian moment invariants are positively correlated with the success of the method. Overall, a quantitative measure of the total importance of shape features is difficult to determine with the unsupervised machine learning method described here, due to the high variability of the results.

4.5.2 Supervised analysis

The support vector machine analysis gives an idea not of the natural groupings of the data but instead address the question of whether there is enough differences between classes in the training data to be able to create a classifier that can effectively identify future points. The data presented in Table 4.4 seems to indicate that this is the case with the feature sets utilized here. This table indicates the maximum percentage that a classifier was capable of achieving based upon five-fold cross validation over the range of C values investigated and shown in 4.10.

The maximum percentage was for the CVNS data set, containing the Cartesian invariants, volume, number of nearest neighbors and volume-normalized second order moments for each cell. Interestingly, it was this data set that scored best for both the 7x and 8x cotyledon datasets. Looking at a comparison of the success of classification by feature set in Figure 4.19, one can see the overall shape of the response to the inclusion of different features does not vary greatly from the 7x to the 8x cotyledon. Overall, it seems to be that moment invariant shape features such as the Cartesian and Zernike invariants can only achieve between 60% and 70% accuracy. As additional features are added, the overall classification power of the model increases, up to a maximum of around 91%.

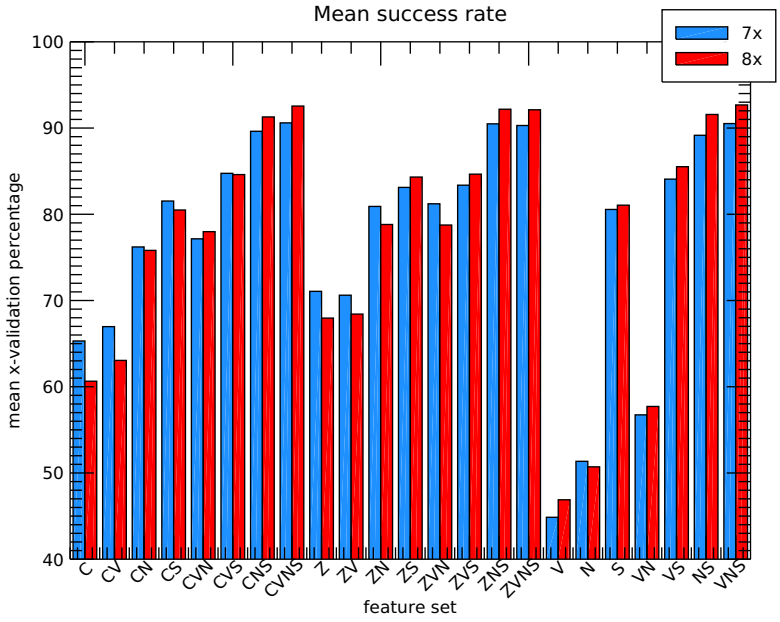

Figure 4.19: Comparison of cross validation success between cotyledons in SVM with indicated feature sets.Values taken from Table 4.4

From Figure 4.19 as well, it can be observed that the volume feature does not appear to play a large role in improving the success of the classifier. Aside from having the worst single-feature classification percentage at between 44-47% for either cotyledon, it also does not improve classification much when combined with other features. When comparing C against CV, CN against CVN, and CVS against CVNS, one may note that the improvements in the percentage success of the classifier are quite low and somewhere on the order of one

to three percent. In fact, this improvement due to the inclusion of the volume feature is reduced to nearly zero with the inclusion of the Zernike features, indicating that there may be some correlation between these two feature sets. However, in the absence of moment invariants, the volume feature alone does not dramatically impact the success rate, with the change from N to VN being approximately 5%, and only about 3% from NS to VNS.

It is also clear from both Figure 4.19 and Table 4.4 that the volume-normalized second order moments (S), which contain orientation and aspect ratio information, are quite successful at generating a good classification model. In fact, based upon the maximum classification percentage by single features, they are the best, with a success rate of between 80-81% for both the 7x and 8x cotyledons.

Nearest neighbors also contributes to the correct identification of cells. While the classification percentage based on nearest neighbor values alone is quite low compared to the shape-based features (Cartesian, Zernike, and second order moments at slightly above 50%), the NN feature set clearly interacts well with others, as can be seen when comparing C against CN, Z against ZN, CVS against CVNS, and ZVS against ZVNS. In these cases, inclusion of the NN set improve classification by between 6-15%.

Examining just the moment invariants, both Cartesian and Zernike performed relatively well alone, at roughly 63% and 70% success rates respectively. These features contribute to the overall success rate when combined with volume or nearest neighbors, as can been seen when comparing V with CV or CVN with VN, but in the presence of the second order invariants, these features seem to contribute little, as can be seen between CNS and NS or CVNS and VNS.

Statistical analysis

Simple numeric comparisons or examinations of bar plots may suggest the relative importance of feature sets, but they do not provide enough information to make any definitive statements. The results of multiple linear regression in Table 4.5 and the results of ANOVA shown in Figures 4.11 and 4.12 give some more insight into what are the important features and interactions between these features and how they contribute to the overall accuracy of the generated SVM classifier.

The histograms of residuals from ANOVA testing for both cotyledons show a much healthier state of affairs than those for the unsupervised analysis in Figures 4.3 and 4.4. Even with the inclusion of fewer data points representing the results of SVM classification, it is clear that the residuals have distributions that are much closer to normal, indicating that there are few if any underlying effects that have yet to be captured by the model.

The ANOVA table for the 7x cotyledon indicates that most feature sets have some kind of an impact on the success of the resulting SVM classifier. In the case of the 7x cotyledon, all listed features, either main effects or their interactions are statistically significant, with the exception of the interaction between Cartesian invariants, volume, and nearest neighbors. In the case of the 8x cotyledon, all main effects and their interactions are significant except the interaction of volume and second order moments, and the interaction of Zernike invariants, nearest neighbors and second order moments.

Every main features has a statistically significant effect on the results of the SVM classification. This indicates that the classes of cells between which a distinction has been drawn do differ with respect to these features, and that these differences are being captured by the classifiers. This is true even with a classifier constructed using the worst performing feature set, volume, which even so, succeeds at a rate almost twice that of random (46% versus 25%), and this being achieved in a one-dimensional data space.

The fact that many interactions are statistically significant is somewhat indicative of what is already known, that each feature set quantifies an aspect of the cell which is not quantified in the other features, although there may be correlation between these features in the given system. As these features are added to and analyzed by the algorithm, they open up additional dimensionality for the SVM algorithm to work in, as well as expose natural groupings or delineations in the data space which can be exploited by the algorithm.

The multiple linear regression table (Table 4.5) for the SVM is utilized, as in the case of the unsupervised analysis, to quantify the effect of each feature set on the final clustering/classifier, and the relative importance of each feature set in the analysis. In this case, the generated model has a correlation coefficient (R_m^2) of 0.81 and 0.85 for the 7x and 8x cotyledons respectively. This is much larger than in the case of the unsupervised analysis, which indicates that the linear SVM model is more appropriate overall in this case. As pre-

viously stated, this value can be interpreted as the fraction of the variation in the overall data that is explained by the linear model [149].

The residual plots in Figure 4.13 are also have much less linear correlation than those of Figure 4.5. However, there is clearly an inverted, possibly parabolic relationship between the predicted success rate and the residuals, indicating that there may be underlying relationships that are not fully captured using the MLR model. The F-value for these MLR models is also quite high at 90.5 and 125.5 for 7x and 8x, respectively, above the critical F-values [144], indicating that the linear model is at least a better predictor than the total overall average success rate [146].

Looking exclusively at Table 4.5, there are positive correlations (R_t) in all features except for volume for both data sets. Using the squared semipartial correlation coefficient for each feature set as a measure of the relative importance of each feature [147], the most important features for the 7x data set are, in descending order, second order moments, Zernike invariants, Cartesian invariants, nearest neighbors, and volume. In the case of the 8x data set, important features ranked in order are second order moments, nearest neighbors, Zernike invariants, Cartesian invariants, and volume. In both cases, nearest neighbors and an invariant feature have very similar values, indicating that each has a similar influence on the SVM success rate.

In the case of the volume feature, the correlation is negative for cotyledon 7x, and in both data sets, has the smallest correlation value. The "coefficient" value gives a measure of the average increase in the SVM success rate with the inclusion of that particular feature. Volume has the lowest value among these, supporting what could have been guessed from the bar plot of Figure 4.19: volume has only a little (but still statistically significant) effect on the overall quality of a generated classification model in the cotyledons examined here.

This fact has some interesting implications with respect to the generalizability of the models developed. If a model has good success regardless of the volume of the cells being quantified, then the magnification of the images used to construct the 3D volume is of little importance, as long as a sufficient magnification is maintained to capture the details needed to obtain measurements of the other features. Any model developed on one cotyledon, if the models do prove to be applicable across individual specimens, would not require an

investigator to do additional work to either match acquisition parameters or transform the data afterwards. Instead, a model could be applied on any data, regardless of the relative sizes of cells.

The confusion matrices shown for the 7x and 8x data sets in Figures 4.15 and 4.16 display actual cell labels, and assigned cell labels for a single SVM classifier generated using the indicated feature sets. It appears that for both data sets, the moment invariant features were most successful in the identification of spongy mesophyll cells, while at the same time being particularly poor in the identification of epidermal cells. There are clear differences between the Zernike and Cartesian invariants as well, which is to be expected, as they quantify shape differently, due to their basis functions. The best shape feature, the second order moments, displayed a relative consistency at identifying all cell classes with the exception of the epidermal cells, which were once again problematic. Ultimately, it requires a combination of features which quantify a variety of aspects of a cell to create a classifier which is similarly effective on all classes.

The supervised analysis makes it clear that all of the features utilized had an effect on the classification success of generated SVM classifiers. It is also clear that their interactions were significant, as the most successful classifiers were generated using several feature sets that quantified a variety of cell information. Finally, the most important features appeared to be the volume-normalized second order moments, which quantifies shape, primarily aspect ratio, and orientation. The second most important features were either the Zernike moment invariants or the nearest neighbors value. The least important feature was volume.

4.5.3 Comparing the approaches

When compared to the MLR analysis of the unsupervised machine learning method, the support vector machine analysis is far more capable of describing the variability of its results through a linear model which lends itself to determining the overall importance of feature sets.

While the Cartesian invariants played a larger and more positive role in the mapping and clustering method, the Zernike invariants tended to contribute more to the success of the SVM method. This is likely due to the fact that the Cartesian invariants contain

redundant information. There are 242 of these values but they only serve to quantify shape up to the 4^{th} order. By way of comparison, the Zernike invariants use fewer dimensions, 121, to describe much more of the shape up to the 20^{th} order. In the unsupervised method, the redundancy and lack of detail of the Cartesian invariants is an advantage. On the other hand, the Zernike invariants may capture minute details of the shape which are not as important for the purpose of making a distinction between classes as more general details that are contained in lower order invariants. However, in the unsupervised method, all dimensions are given equal weight and importance; meaning that these small details can mask the general shape. This is not a problem with the Cartesian invariants.

The support vector machine technique treats the high-dimensional space differently than t-SNE. Instead of giving each dimension equal weight, it identifies hyperplanes which optimally separate the data. If a dimension is not helpful in creating a classifier, the generated hyperplane will be parallel to that dimension, which prevents unhelpful data from confusing the classifier.

For both methods, the volume-normalized second order moments, which contain aspect ratio information and orientation information in six values proved to be the most important feature. While its positive influence in cell classification is not altogether unexpected, the degree to which it outperformed the other features is remarkable. It is probable that the reason orientation information is so powerful is that these cotyledons have evolved to exist in a 3D environment in which orientation is an important factor. Palisade mesophyll must support the bottom of the leaf, and must be oriented vertically to provide optimum mechanical support. The top of the epidermis must be turned upwards to expose the spongy mesophyll to the sunlight and allow for easier gas exchange. While the second order moments do contain shape information, the invariants, as quantifiers of pure shape, do not contribute to the success of the classifier as much as orientation does. This is not to say that *any* single feature is enough. It takes a variety of information to generate a classifier that is over 90% successful.

Utilizing purely local features, those which are properties of the cell itself and divorced from the overall cotyledon, classifiers can be constructed that are capable of achieving approximately 70% success (utilizing SVM and the ZV feature set). A very popular method

of cell cytometry, flow cytometry, flows cells in solution through a channel, one at a time, to allow measurement of the cell. Recently, a method of 3D shape reconstruction has been developed that would allow the reconstruction of a cell's shape based upon wavefront analysis of the fluid in the channel generated by the random rolling of the cell [119]. If a more successful SVM classifier could be constructed using moment invariants and size as described here, one could perform automated cell identification based on shape in flow cytometry.

4.5.4 Generalized classifiers

Table 4.6 summarizes the results of training a classifier on one cotyledon and applying it to another. The method was quite successful, with the best classifiers achieving over 90% success. The best feature sets worked both ways, from 7x to 8x and from 8x to 7x, and were ZNS, VNS, and ZVNS. It is interesting to note that when training a model on the 7x data with Cartesian invariants, it applies less well to the 8x than one trained on the 8x data with Cartesian invariants applied to the 7x model.

The success of this generalization of the SVM classifiers indicates that it may be possible to construct a general SVM classifier for the species. If this is possible in cotyledons of *Arabidopsis Thaliana*, then it could potentially be extended to other systems where automated identification in 3D is useful, such as human blood or other tissue, allowing for identification of irregular or diseased cells.

4.6 Conclusions

This chapter has shown to what extent shape, as quantified through moment invariants, plays a role in the identification of cell function in cotyledons of *A. Thaliana*. Utilizing unsupervised machine learning methods, it was shown that there are natural delineations in the data which allow the clustering of vascular, epidermal, spongy mesophyll, and palisade mesophyll cells together without any identification from a user. The success of this method was shown to be due primarily to orientation information stored in volume-normalized second order moments, and shape information in the Cartesian moment invariants. However, due to the stochastic nature of some algorithms, and the sensitivity of the success of this method

to input parameters, it is difficult to draw quantitative conclusions from the unsupervised machine learning data.

Using a support vector machine classifier trained with user-identified cells, it was possible to classify cells correctly with a success rate of up to 92%. Multiple linear correlation showed that Cartesian and Zernike invariants were responsible for between six and fifteen percent of the total variation in the success of these classifiers, with the most important features being the second order moments at 60-70% of the variance. With all other features held constant, the success of these classifiers improved between eight and twelve percent with the inclusion of moment invariants. No one feature could achieve the best success rate, which utilized a variety of features including those pointed out earlier and nearest neighbors and volume. Finally, it was shown that SVM classifiers could be applied across specimens, indicating that a general classifier for the cotyledons of the species may be possible. Therefore, the hypothesis is quantitatively supported by the SVM analysis of this chapter.

CHAPTER V

Shape-based classification of metal powders for additive manufacturing

5.1 Background

5.1.1 Additive manufacturing

Additive manufacturing (AM), the fabrication of a three-dimensional part through the addition of layers of material, usually done by a computer driven mechanism based on a CAD file, has been of interest to researchers since the 1980s, although its roots go back far further [150]. Early methods of AM (also known as 3D printing and, erroneously, rapid prototyping, among other names) were based on photopolymers [151, 152]. Approaches utilizing polymer powders and selective laser sintering (SLS) quickly followed [153–155]. A distinction should be made between selective laser sintering and selective laser melting (SLM). As the name suggests, with SLM, the material is fully melted, creating a part of higher density with less post-processing required.

The additive manufacturing of metal components has been an area of interest for researchers since the early 1990s [156–159] although patents exist from earlier [160]. These early methods, like those that are used today, utilize metal powders as the feedstock, and a high-powered laser as the heat source.

There are enormous and obvious advantages to additive manufacturing that set it apart from more traditional methods that require the removal of material. The first is that a nearly or fully finished part can be fabricated very quickly, needing minimal post-fabrication machining to finalize the process. Second, due to the layer-by-layer nature of the construction, it is possible to construct parts with features that would be impossible with conventional

machining processes. Finally, the overall process is, by its nature, much more efficient with material and generates less waste than other methods.

There are drawbacks to the SLM and SLS techniques. The speed at which it is possible to fabricate a component, along with the nature of the laser used as heat source, also means that the material is subjected to extreme temperature gradients, as well as fluctuations and inhomogeneities in the distribution of heat, which lead to the majority of problems with the method. The rapid heating and cooling leads to tension and compression within the material, the overall effect of which is a buildup of a large degree of internal stresses within a part. In some cases, the stress can meet or exceed the yield stress of the material [161–163]. The rapid melting and solidification of the material, along with the powder nature of the feedstock also leads to another primary problem: porosity. This can dramatically influence the mechanical properties of the finished part [1, 164, 165]. Further defects include oxide inclusions, unmelted particles, and cracking [166]

There are a large number of parameters that play a role in influencing the formation of the above defects, and in affecting the bulk mechanical properties of the final part as a whole. These are illustrated in Table 5.1. Ultimately, the purpose of this chapter is to prove the usefulness of moment invariants as a metric by which the powder-related property of particle shape can be more fully understood. A number of published works have noted a correlation between the shape distributions of powders and the final properties of the bulk material (as discussed in Section 5.1.2), but without a method of shape quantification that can completely capture the details of a full shape, this connection cannot be fully explored.

Laser-related	**Scan-related**	**Powder-related**	**Temperature-related**
Laser power	Scan speed	Particle size	Powder bed temperature
Spot size	Scan spacing	Particle shape & distribution	Powder feeder temperature
Pulse duration	Scan pattern	Powder bed density	Temperature uniformity
Pulse frequency		Layer thickness	
		Material properties	

Table 5.1: Parameters of SLS/SLM that affect the final properties and the presence of defects in a final part. Modified from [1].

5.1.2 Particle shape and material properties and defects

In the selective laser sintering and selective laser melting processes, the primary interaction is between the laser and the powder to be melted. In this system it is the properties of the feedstock that control what the optimal laser-related, scan-related and temperature-related parameters are. In particular, bulk material properties, such has the melting temperature and thermal conductivity, control the flow of heat through the powder and allow for the particles to achieve melting temperature. However, the particle nature of the material means that heat can be transferred either radially through the gaps between individual grains of powder, or conductively where the particles are touching. This means that the density of the powder, which is closely tied to the morphology, plays a large role in determining the way heat moves through the material [167, 168].

In a study of particle shape on the properties of a Ti-TiB composite material produced by SLM, Attar *et al.* noted a significant difference in density and performance of two samples fabricated with powder of differing morphologies [169]. A specimen constructed with particles of "irregular" morphology displayed a relative density of 95.1%, and an ultimate compression strength and strain of 883 MPa and 5.5% respectively. Another sample built using spherical particles had a density of 99.5%, and nearly twice the compressive strength at 1421 MPa. The authors suggested that the improvement was due to the improved flowability and packing density of spherical particles.

Similar results were found by Niu *et al.* in the selective laser sintering of high speed steel (HSS) powders [170]. Four particle classes were examined: gas atomized of roughly spherical shape, and water atomized of more irregular morphology. It was found that the gas atomized powders with more regular morphologies generated more homogeneous microstructures overall, while the irregular water atomized powders formed agglomerates and porous surfaces.

Olakanmi found similar results in aluminum powders [171], with more irregular particulate shapes leading to an increase in agglomerates and the formation of pores. Spherical particles tended to correlate with more dense and homogenous regions.

One of the major advantages to powder bed additive manufacturing, in which a layer

of powder is deposited in the working area, is selectively bonded by the laser, and then another layer of powder applied, is that the powder that is unused to create a bulk part can be reused. Tang *et al.* found that in the case of Ti-6Al-4V powders, the oxygen content and the morphology of the feedstock slowly changes with repeated reuses [172]. It was also found that particles became less spherical with increasing reuse times, and that fewer satellite particles (small attached particles) were observed as well. The authors also noted that flowability increased, which, although possibly reduced by the reduction in sphericity, was most likely influenced by the reduction in satellite particles.

5.1.3 Powder characterization

While some attention has been paid to the impact of powder morphology on the properties of a manufactured material, as listed above, the majority of these studies have failed to quantify this morphology in any way, but instead resort to qualitative analyses. Recently, more attention has been focused on quantifying the shapes of additive manufacturing feedstock using a variety of methods. Some of these are summarized here.

Slotwinski *et al.* utilized a wide variety of methods to characterize nearly all aspects of commercial stainless steel and cobalt chrome powders, in order to obtain quantitative measures of variance between batches of nominally identical feedstock [173] . Following previous work on concrete particles [174], Slotwinski *et al.* matched three-dimensional reconstructions of individual particles, obtained through X-ray computed tomography, to spherical harmonics series. The coefficients of the harmonics were stored and used as shape descriptors. In the end, these spherical harmonics series coefficients were used to calculate more basic shape information, such as volume, surface area, width, thickness, and even the moment of inertia tensor, coming very close to the shape descriptors used in this work. It was found that there was little to no variation between samples.

As an aside, there has been a significant amount of work in the field of morphological analysis as applied to aggregates used in concrete. A variety of shape quantifiers have been used, including the moment of inertia tensor (a second order moment tensor, previously mentioned in Section 3.1.3 [175]). However, recently, the focus has shifted toward using spherical harmonics [176–178]. The spherical harmonic approach is similar to that of 3D

moment invariants, but suffers from the fact that the shapes that are quantified are limited to those to which one can draw a line from the center of mass to every point on the surface without exiting the object, referred to as "star shaped".

Strondl *et al.* quantified morphology through the use of a commercial system (Camsizer XT; Retsch Technology GmbH, Haan, Germany). The features measured were the size and the aspect ratio of particles. No significant difference was found between fresh and recycled powder via these methods. Correlations between size and morphology to particle flow rates were noted but correlations to final build properties were negligible. Markusson [179] utilized a shape factor value, related to the particle perimeter to area-equivalent circle perimeter. Some correlation was found between shape factor and flowability for Inconel 718 powder. Tan *et al.* noted the prevalence of simple aspect ratio and Feret diameter measurements in the field [180].

An alternative method of powder analysis based on computer vision and machine learning techniques was conducted by DeCost *et al.* [181]. For eight different feedstock powders, approximately twenty backscatter micrographs were taken from a random sample of each. Corner and blob interest points were identified, and scale-invariant feature transform (SIFT, [182]) was applied at each to generate a vector of 128 values describing that point. Some of these values were then clustered into 32 groups using k-means clustering. The distance from each SIFT value was calculated to the centroid of its cluster, generating a 4096 dimensional vector, reduced to 32 dimensions using PCA. From here, an SVM classifier is created which classifies images as being of one type of powder. The resulting classification was over 95% effective, indicating that important and distinguishing features of the images were being quantified.

The purpose of this chapter is to split the difference between the more simple and human-selected measures described above, such as aspect ratio, and the machine learning methods of DeCost *et al.*. The advantages of the quantities that are selected by hand is that their relation to physical shape is more direct and intuitive, while the machine learning methods have the advantage of being selected by an algorithm in an automated fashion and in a way that optimizes the capture of important aspects of the feedstock. The use of moment invariants, which can describe the morphology in a way that is more related in a

one-to-one way with the shape of the powder, combined with machine learning methods, can be used to identify and quantify powders comprehensively and in a way that can, in the future, be related to powder properties. This work shows how moment invariants can quantify important features of the types of shapes displayed by these particles and can make meaningful distinctions between these classes of shapes.

5.2 Materials and methods

Data sets

The backscatter SEM micrographs used in this chapter and which were the primary data sets, were provided by, along with the advice and help of, Dr. Brian DeCost. Details of their collection are given in [181], but are summarized for completeness.

Data sets consisted of backscatter SEM micrographs of eight separate types of additive manufacturing feedstock powder. Of these eight samples, five were designed for use with an EOS machine (EOS GmbH, Krailling, Germany). These materials (and their shortened identifiers) are AlSi10Mg (Al-EOS), Inconel 718 (In-718), maraging steel (Mar-steel), type 316 stainless steel (SS-316), and Ti-6Al-4V (Ti64-EOS). The remaining powders were designed for use with an Arcam machine (Arcam AB, Mölndal, Sweden), however, two of these were manufactured by other companies (AMETEK Inc., Berwyn, PA), (Starmet Corporation, Concord, MA). These materials are both Ti-6Al-4V (Ti64-Ametek), (Ti64-Starmet).

These powders were mixed to ensure an even distribution of particle sizes, and then blown across carbon adhesive tabs on SEM stubs using compressed air. Images were taken on an SEM using backscattered electrons, and attempts were made to compensate for differences in density and particle size by adjusting the magnification. Between 23 and 41 images were taken of each powder, and care was taken to ensure that no two images overlapped, including key-point matching with RANSAC filtering, which is capable of detecting regions of images that are related to each other by translation [183].

Examples of these images are given in 5.1. It may be noticed that no scale bars are present. Because moment invariants are invariant to size, the magnification of the image is irrelevant to the analysis. This has the advantage of not requiring any sort of preprocessing in order to conduct this analysis.

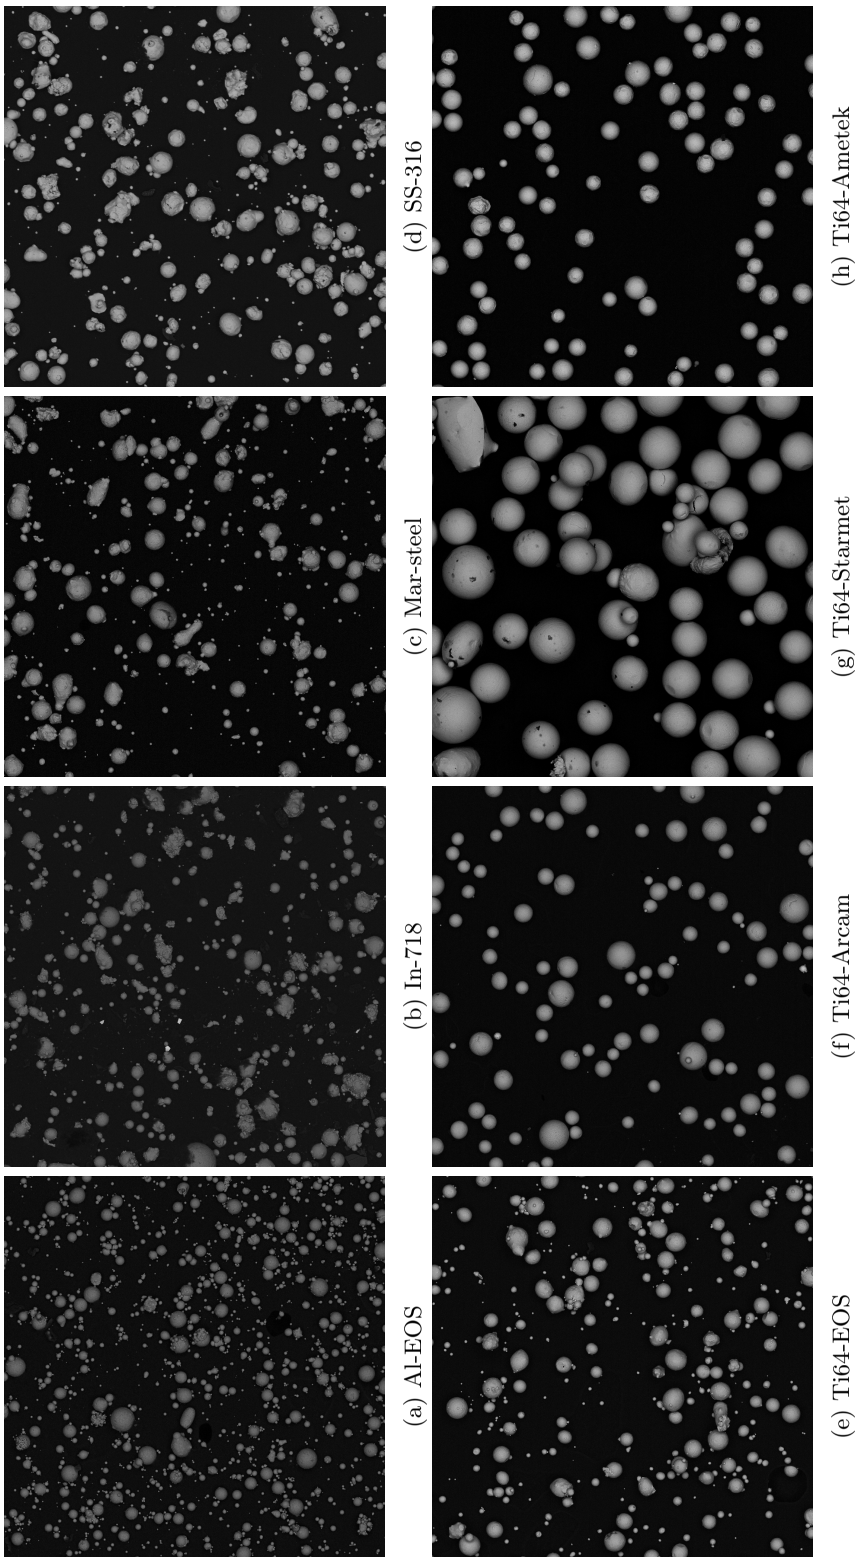

(a) Al-EOS (b) In-718 (c) Mar-steel (d) SS-316

(e) Ti64-EOS (f) Ti64-Arcam (g) Ti64-Starmet (h) Ti64-Ametek

Figure 5.1: Examples of powder images used in this analysis. Courtesy of the Holm Group, Carnegie Mellon University. Degree of magnification and therefore particle size was not included in this analysis.

5.2.1 Segmentation and feature extraction

The methodology for image analysis and the extraction of shape information used here is similar to that of Chapter IV, but in two dimensions. First, it is necessary to segment each image. For this, as with the cross-sections of the cotyledons, *ilastik* was utilized due to its excellent flexibility and robustness to variations in the images. For each material, a classifier was trained, using selected subregions from several images. These classifiers were then applied to the entire set of images to label the background (values of zero) from the individual powder particles (values of one). This portion of the data analysis may be the most important, due to the nature of the shape feature extraction. In order to quantify the shape of an individual particle, the moment invariant calculating algorithm looks at regions of adjoining pixels with a value of one. If even a single pixel is out of place and connects two nearby particles which should not be connected, the new agglomerated region will be erroneously considered a single shape and quantified as such.

From here, the central moments of identified regions were calculated according to the method of Novotni and Klein [10], and were subsequently used to calculate the two ω and three τ invariants as described in Section 3.1.2. In addition, using equations developed by Suk and Flusser [106], the eighty affine Cartesian invariants up to the 12^{th} order were calculated for each particle as well. These sets were stored in data arrays which contained material identity, image identifiers, and a vector of shape values for each particle.

5.2.2 Data analysis

In order to explore the ability of moment invariants to make meaningful distinctions between the classes of shapes displayed by AM powder feedstock, it is necessary to use a method of analysis that can quantify differences between distributions. In this work, each particle is quantified separately, and an image of the powder is treated as a distribution of the properties of this powder. If the shape of each individual particle can be reduced to either five (ω/τ invariants) or eighty (affine Cartesian invariants) continuous values, then the morphology of each particle can be thought of as a point in a five- or eighty-dimensional shape-space. Any collection of particles and any sample of powder therefore can be represented as a distribution within that shape-space. Drawing conclusions regarding powder

morphology, be it the similarity, degree of difference, or the identity, for a sample of powder is then a matter of comparing distributions.

Several methods were utilized in this work to attempt to quantify the differences between the shape distributions of individual powder samples. One of the problems encountered early on was the way in which the dimensionality of the space directly impeded the ability to relate distributions. If one treats the space as a histogram, even a five-dimensional histogram with 100 bins in each dimension has 10^{10} individual bins. If one reduces the number of bins to attempt to correct for this, even a reduction to 10 bins leads to 10^5 bins, and this comes with a dramatic decrease in the degree of discriminatory power the analysis is capable of. With 259 images, it is easy to see how sparse even a small histogram of this dimensionality will be.

The first attempt to get around this problem was the modification of previously utilized methods that have already been shown to be capable of the identification of 3D objects by the analysis of the shapes of random 2D cross-sections through these objects [9]. In this approach, second order moment invariant maps (SOMIMs) and projected moment invariant maps (PMIMs) are generated for each image, the details of which are described in Section 3.1.2 . Both maps are essentially a 2D histogram of a shape distribution. In the case of the SOMIM, the x and y axes are the values of the ω_1 and ω_2 invariants respectively. As the name suggests, the PMIM is a projection from a five-dimensional space defined by the five invariants: $\{\omega_1, \omega_1, \tau_1, \tau_2, \tau_3\}$. This projection places the average ω value on the x axis and the average τ value on the y.

In order to determine from which material each image originated, the Hellinger distances between a test image's SOMIM and PMIM were calculated. The material to which the Hellinger distance was shortest was reported as the identified material. The Hellinger distance between two PDFs is given by:

$$H(p,q) = \frac{\sqrt{1-\beta(p,q)}}{\sqrt{2}} \tag{5.1}$$

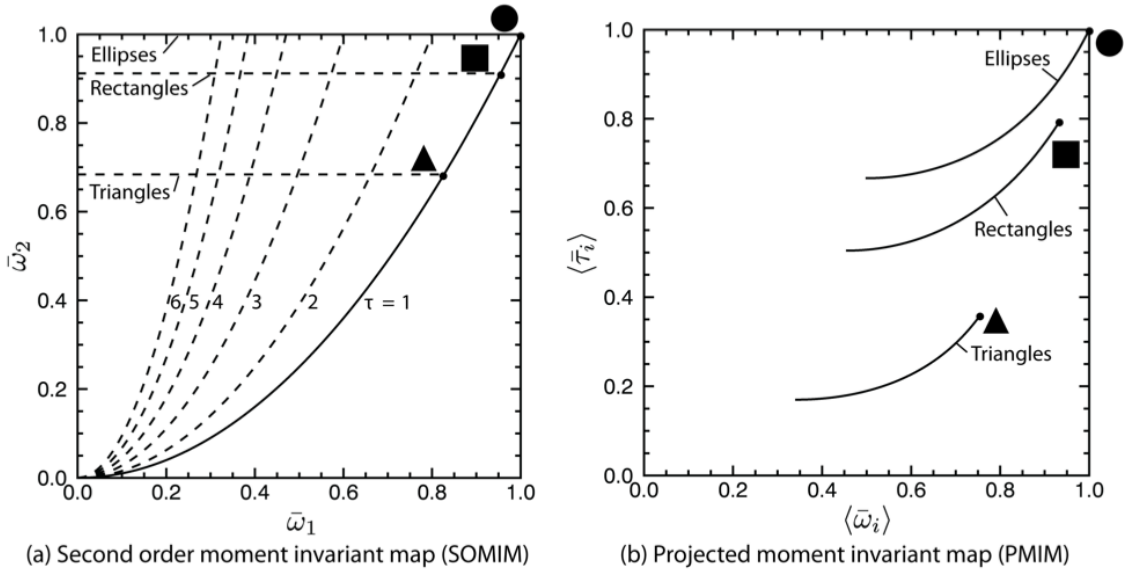

Figure 5.2: Example moment invariant maps with identified shape regions.

where $\beta(p, q)$ is the Bhattacharyya coefficient, given by:

$$\beta(p,q) = \sum_{i=1}^{N} \sqrt{p(i)q(i)} \quad \text{with} \quad \sum_{i=1}^{N} p(i) = \sum_{i=1}^{N} q(i) = 1 \tag{5.2}$$

While the Bhattacharyya coefficient is not a metric, the Hellinger distance is. A metric utilizing the combined distance between both two SOMIMs and two PMIMs was also used, following [9], the metric being: $\sqrt{H_{\mathrm{S}}(p,q)^2 + H_{\mathrm{P}}(p,q)^2}$, with H_{S} being the Hellinger distance between SOMIMs and H_{P} the Hellinger distance between PMIMs.

Another attempt at circumventing the dimensionality problem was directly inspired by DeCost *et al.* [181] and the use of a bag-of-words approach [184]. In this method, the values of a distribution are not directly quantified by generating a histogram of the quantifying feature, but instead clustered to identify (in this case) particles that are similar. Because the clustering is done in a moment invariant space, the clusters should identify particles which have commonalities of morphology. As seen in Chapter IV, the clustering of these shape features can lead to similar shaped objects being placed together.

To use AM powder as an example, each image is represented as a distribution of points representing particles in an m-dimensional shape space. These points are then clustered, each cluster representing some underlying feature or collection of features which are shared

among particles. With the bag-of-words approach, these clusters, or words, are now used as the identifying features, rather than the underlying moment invariant values. If there are n words, then an image, instead of being represented as an m-dimensional histogram, can be represented as a one-dimensional histogram with n number of bins. This method both works to reduce dimensionality, and ensures that the features being measured are relevant to the question at hand.

Clustering here was performed using t-SNE and HDBSCAN, as described previously in Chapter IV. This method, while not generating a predefined number of clusters, has advantages over a simple clustering technique such as the k-means algorithm. First of all, utilizing Cartesian moments generates an 80 dimensional space, which is quite large and leads to problems related to the curse of dimensionality (mentioned in Section 4.2), meaning that some kind of dimensionality reduction is necessary prior to clustering. t-SNE is advantageous when dealing with moment invariants, as these values tend to be very continuous, and show a large degree of variation in a large number of directions. t-SNE manages to maintain a large degree of the structure of the data, while a more common and simple method, principle component analysis, does not.

To illustrate this point, Figure 5.3 displays two mappings of the entire set of points examined in this work. The first is a mapping using PCA, and the second is one which utilized t-SNE. Identifying natural clusters from the PCA plot is nearly impossible, while there is clear structure in the t-SNE mapping. The PCA mapping likely failed due to extreme outliers. These particles displayed extreme morphologies which led to large moment invariant values in two dimensions. It is probable that the differences in the morphologies between these particles is subtle, and during PCA, the small variations which matter when observing particles with the eye, such as number of satellite particles or subtle differences in the roundness of the object, get washed out in the large variations between a particle that is a large agglomerate of several smaller spheres and the more common rounder particles. t-SNE is so successful at this mapping, probably because it is not so much concerned with variation, but instead with retaining the proximity of points as the dimensionality is reduced.

t-SNE is therefore performed on the particles from the remaining images not selected as testing samples. This generates a map similar to that in Figure 5.3. A perplexity value of

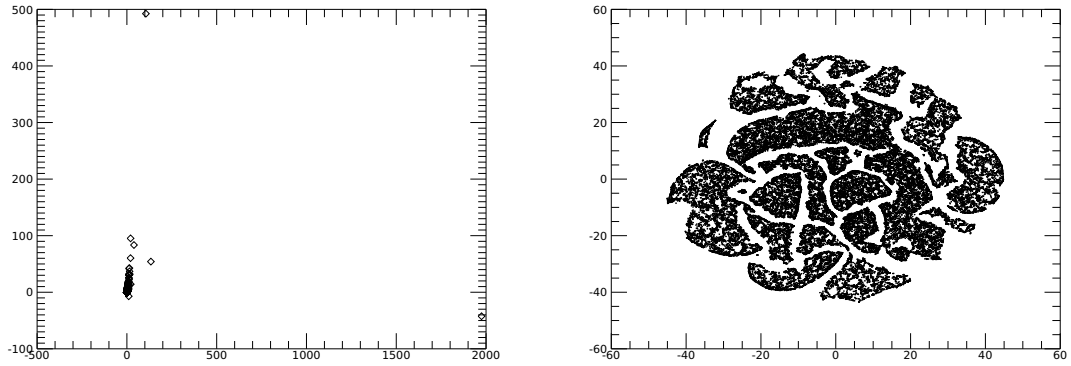

(a) PCA mapping

(b) t-SNE mapping

Figure 5.3: PCA versus t-SNE mapping, from 80-dimensional Cartesian moment invariant space to two-dimensions, of the over 50,000 particles analyzed in this chapter.

50 was determined through trial and error. HDBSCAN is utilized to identify these clusters in an automated way. HDBSCAN allows for the number of clusters naturally present to be determined, rather than forcing a number not necessarily supported by the structure on it. In addition, k-means favors round clusters, which are not necessarily always present in a t-SNE map. The clusters generated by HDBSCAN are the words to be used in the bag-of-words representation.

Both individual training images, and the materials as a whole can now be represented as a histogram of the clusters to which its particles are assigned, which allows matching to the testing images. However, the testing images do not have a bag-of-words representation, and because t-SNE is a stochastic algorithm, the process cannot be repeated to cluster the testing particles. Instead, each particle is related to the particles around it using a k-nearest neighbors (kNN) classification: the k nearest particles in shape-space are used to identify to which cluster the new particle belongs, by simple voting. In this work, a k value of five was chosen arbitrarily. In addition, a distance-based voting scheme was also employed simultaneously and the results compared. In this method, the vote of each neighbor particle is weighted by the inverse of the distance to the particle to be classified.

Using this k-NN classifier, it is possible to generate a bag-of-words representation for every test image. These histograms can then be measured for similarity against the material histograms or the training image histograms via the Hellinger distance. The material of the

image with the smallest distance, or the overall histogram belonging to the material with the smallest distance, being assigned to the test image.

In this work, for each material, twenty random images were chosen and each of the previously described methods were tested over twenty iterations. In each iteration, one image per material was selected as a test image, and the remaining images were used as a training set for each classifier. The resulting classification was saved and compiled into confusion matrices, which are included below. The choice of twenty images per material was an attempt to allow for balanced classification. The smallest number of images for a material is 24 for the maraging steel. By using twenty, nearly the entire set of images of maraging steel are analyzed. Overall there are 267 images, which means that approximately 30% of the images were tested in the above manner.

5.3 Results

5.3.1 SOMIM and PMIM matching

Second order and projected moment invariant maps (SOMIMs) for each material are shown in Figure 5.4. These illustrate the differences between particle morphologies on the level of detail capable of being quantified by second order moments, which are essentially the amount of detail in the equivalent ellipse of each particle. The most obvious commonality between these distributions is the location of the peak. Referring to the SOMIM map in Section 3.1.2, the peaks of these distributions all lie near the point that corresponds to the value of the circle. This is to be expected based upon the micrographs of Figure 5.1. It is clear that the majority of the shapes are circular or close to it.

Of the non-titanium materials, Al-EOS shows the greatest degree of irregularity. For the titanium-based alloys, the powder intended for the EOS machine has a greater spread of shapes. The remaining Ti64 materials, intended for use with an Arcam device, show much more concentrated peaks, indicating little variation in their shapes.

The PMIMs of Figure 5.5 emphasize the spread of morphologies in these samples. The fourth order τ moments, the mean of which is plotted on the y axis, are much more sensitive to changes in shape at the exterior regions of an object. This is because with increasing p and q values, the Cartesian basis function, $x^p y^q$, grows much more quickly as one moves

from the centroid of the object. The PMIMs therefore paint a clearer picture regarding the variety of shapes present in these materials. It is made even more obvious that the Al-EOS samples has the most extreme variation in the shape of its particles. The Inconel, maraging steel, and stainless steel all shown similar degrees of variation, with subtle differences: the peak in the case of maraging steel is more spread out, for example. In the case of the Ti64 alloys, the PMIM magnifies the variation of shapes in the EOS powder. The remaining titanium powders show a high level of similarity, with very tight peaks.

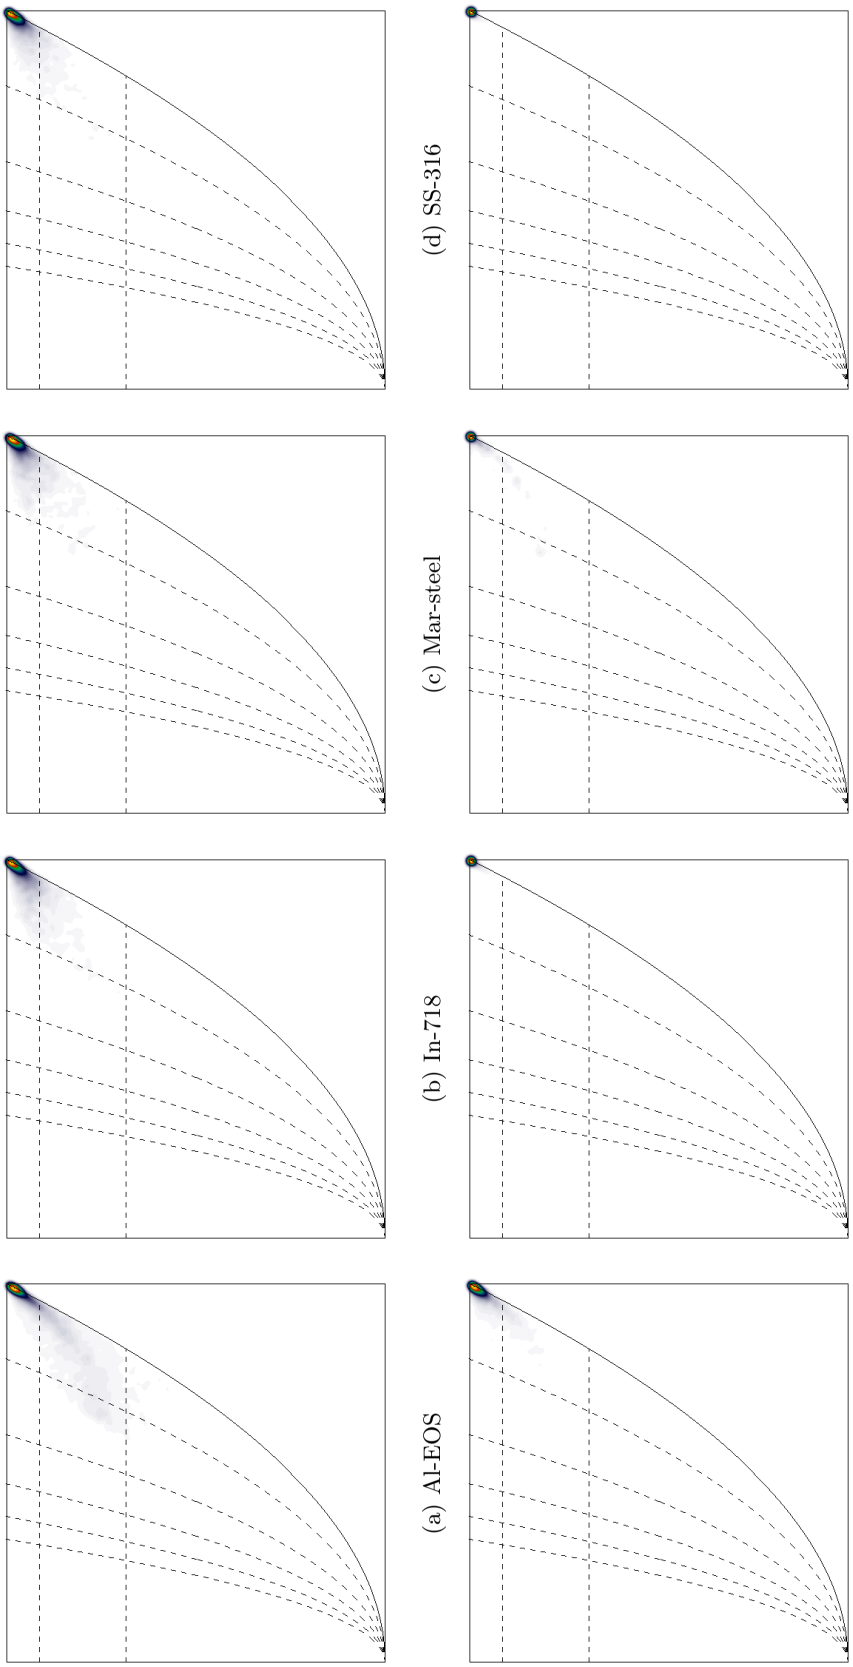

(a) Al-EOS (b) In-718 (c) Mar-steel (d) SS-316

(e) Ti64-EOS (f) Ti64-Arcam (g) Ti64-Starmet (h) Ti64-Ametek

Figure 5.4: Master SOMIMs for each class of powder material.

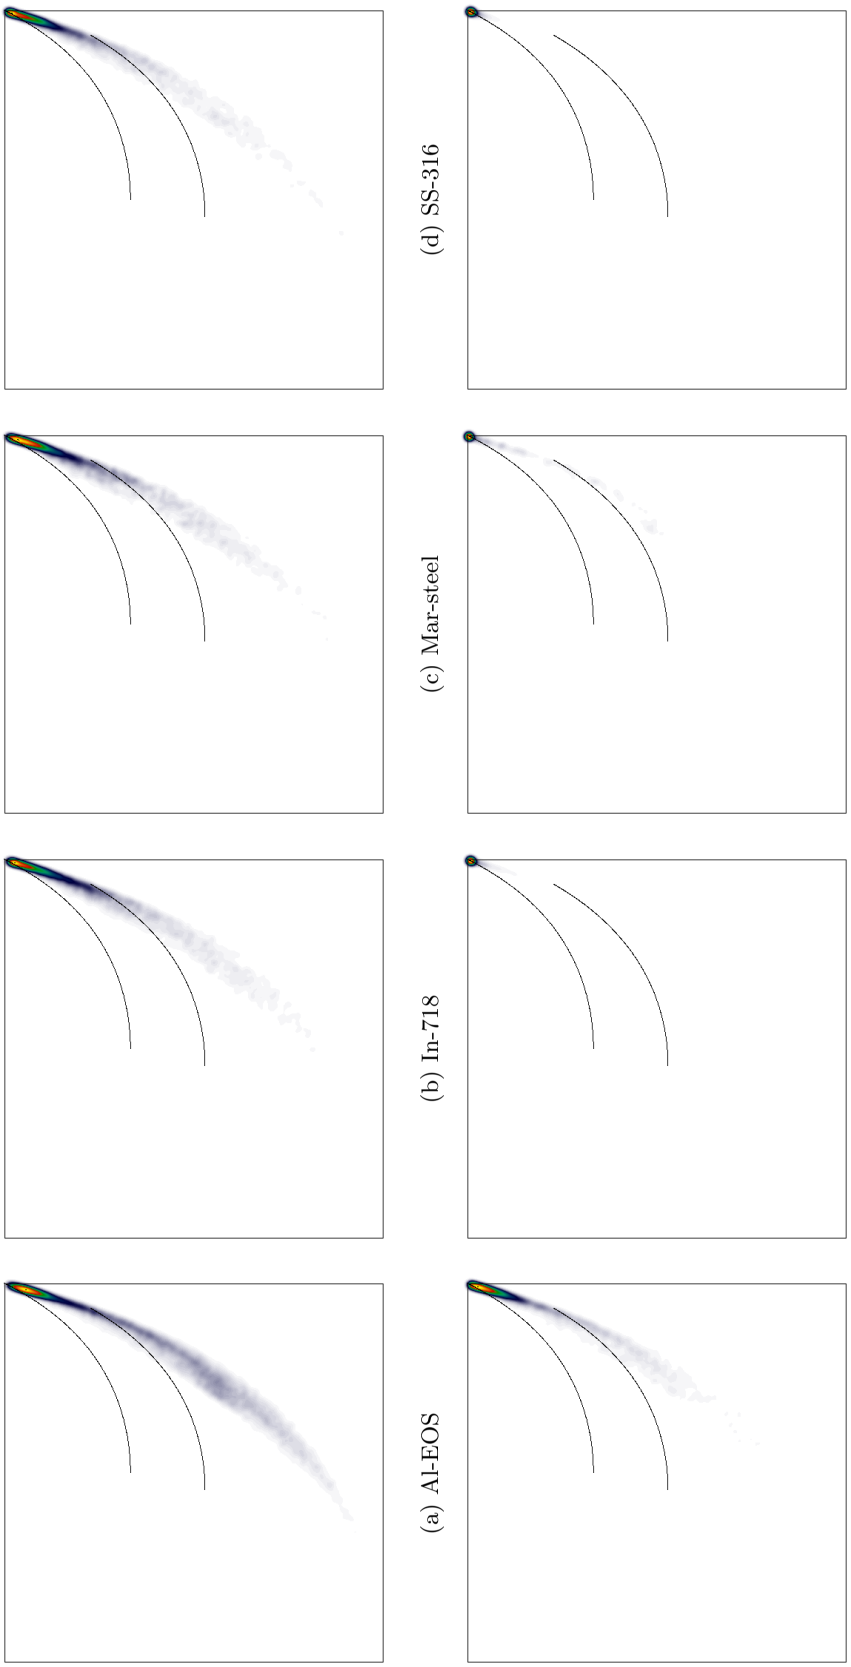

(a) Al-EOS (b) In-718 (c) Mar-steel (d) SS-316

(e) Ti64-EOS (f) Ti64-Arcam (g) Ti64-Starmet (h) Ti64-Ametek

Figure 5.5: Master PMIMs for each class of powder material.

Confusion matrices for the SOMIM, PMIM, and the combined metric are shown in Figures 5.6, 5.7, and 5.8. The SOMIM matching method displayed an overall success rate of 54.37%, better than randomly guessing. This near-half success rate is matched by the success on the materials: in general, only half of them could be matched with a high degree of success. The materials which were most correctly identified were Al-EOS, Mar-steel, Ti64-Ametek, and Ti64-EOS. The method seems to have completely failed in the case of In-718, SS-316, Ti64-Arcam, and Ti64-Starmet.

The success rate of the PMIM-based matching is slightly higher at 55.88%. However, this increased success appears to have been achieved by an overall reduced success of the maraging steel classification, and slightly improved success rates for In-718 and SS-316. So if success is measured by the number of materials for which there is a large success rate, the PMIM-based approach may actually be considered less effective than the SOMIM based method. Finally, the confusion matrix for the combined metric shows success and failure in similar materials to that of the SOMIM metric.

5.3.2 Bag-of-words

The bag-of-words model, as described previously, takes clustered data and treats the distribution of points within those clusters as features. In order to achieve this in the case of these powder micrographs, all of the particles in the training data were mapped to two dimensions using t-SNE, and subsequently clustered using HDBSCAN. An example of this mapping and clustering is shown in Figure 5.9 and Figure 5.10. This also illustrates how this mapping and clustering for the ω/τ invariant space has grouped similar particles together, by placing the particles corresponding to the points on the t-SNE map in their proper location, colored by their cluster.

Figure 5.10a and 5.10b illustrate how similar features tend to be grouped together. The green cluster in the first image clearly consists of nearly perfect circles. The teal particles have large aspect ratios but are otherwise close to elliptical. The blue particles are agglomerates of particles, they have either several small satellites, or consist of several joined particles. In the second image, another portion of the blue region is visible, with similar shapes to particles in the blue region of the first. The red regions displays particles

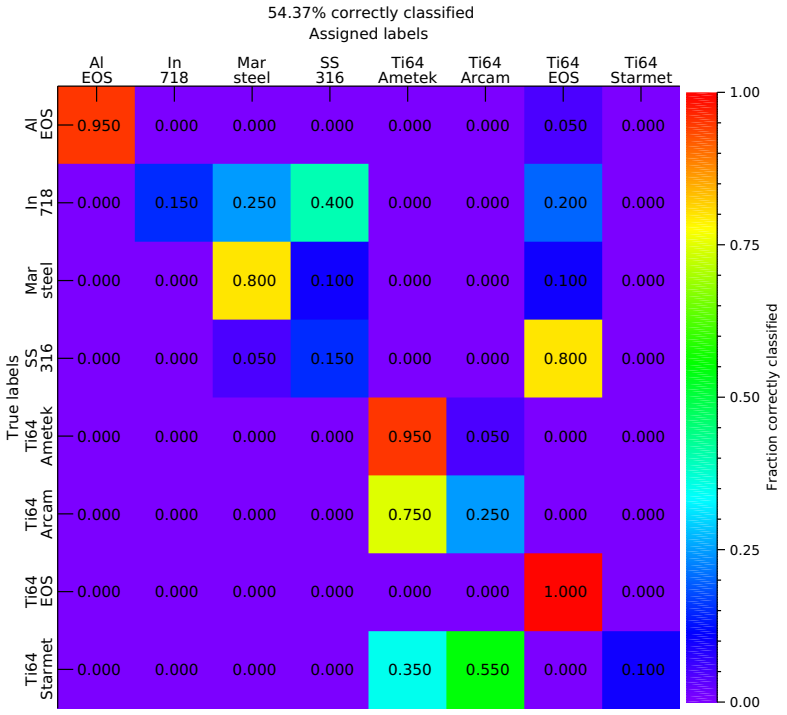

Figure 5.6: Confusion matrix for Hellinger distance mapping of test image SOMIMs to overall material SOMIMs

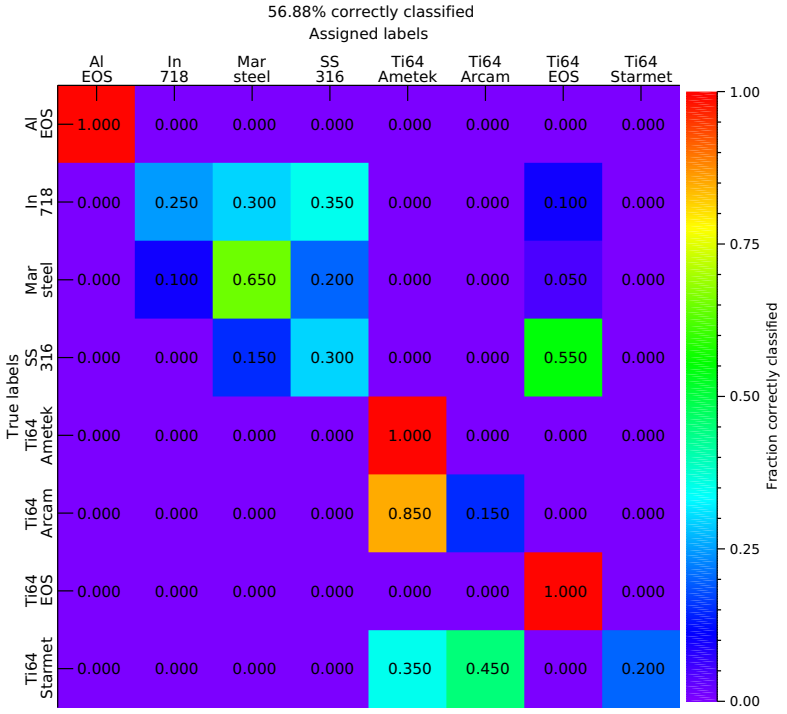

Figure 5.7: Confusion matrix for Hellinger distance mapping of test image PMIMs to overall material PMIMs

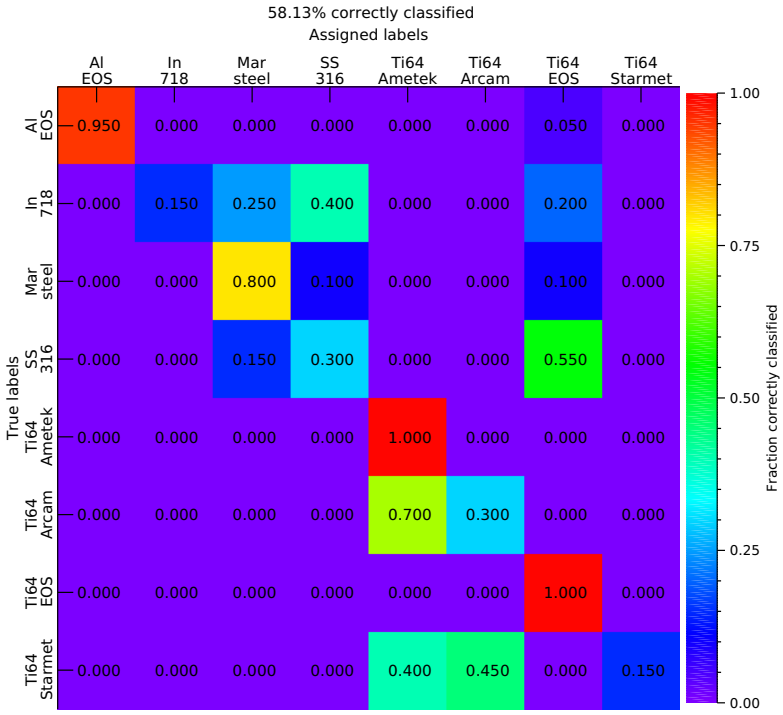

Figure 5.8: Confusion matrix for Hellinger distance mapping of both SOMIM and PMIMs

that are close to circular with some subtle deviations. The last figure, Figure 5.10c shows particles from a single cluster. However, the methodology has apparently failed, as some near circular particles are included along with ones which show extreme deviations from circular.

Figure 5.11 shows the result of a t-SNE mapping from the eighty-dimensional affine Cartesian invariant space to two dimensions for the training set of particles. Three regions are shown, and the particles related to these points in Figure 5.12a through Figure 5.12c. The t-SNE mapping shows somewhat different results than the ω-space mapping. Primarily, there are a number of separate clusters with circular particles. The red, pink, orange and yellow clusters all contain similarly shaped particles. However, the clusters are more coherent than in the ω-based mapping. Clusters show clear similarities. The green cluster in Figure 5.12a shown small numbers of satellite particles, moving the shapes away from circular, but in a relatively small way. The light-blue particles in Figures 5.12a and 5.12b have more extreme morphologies indicative of joined particles.

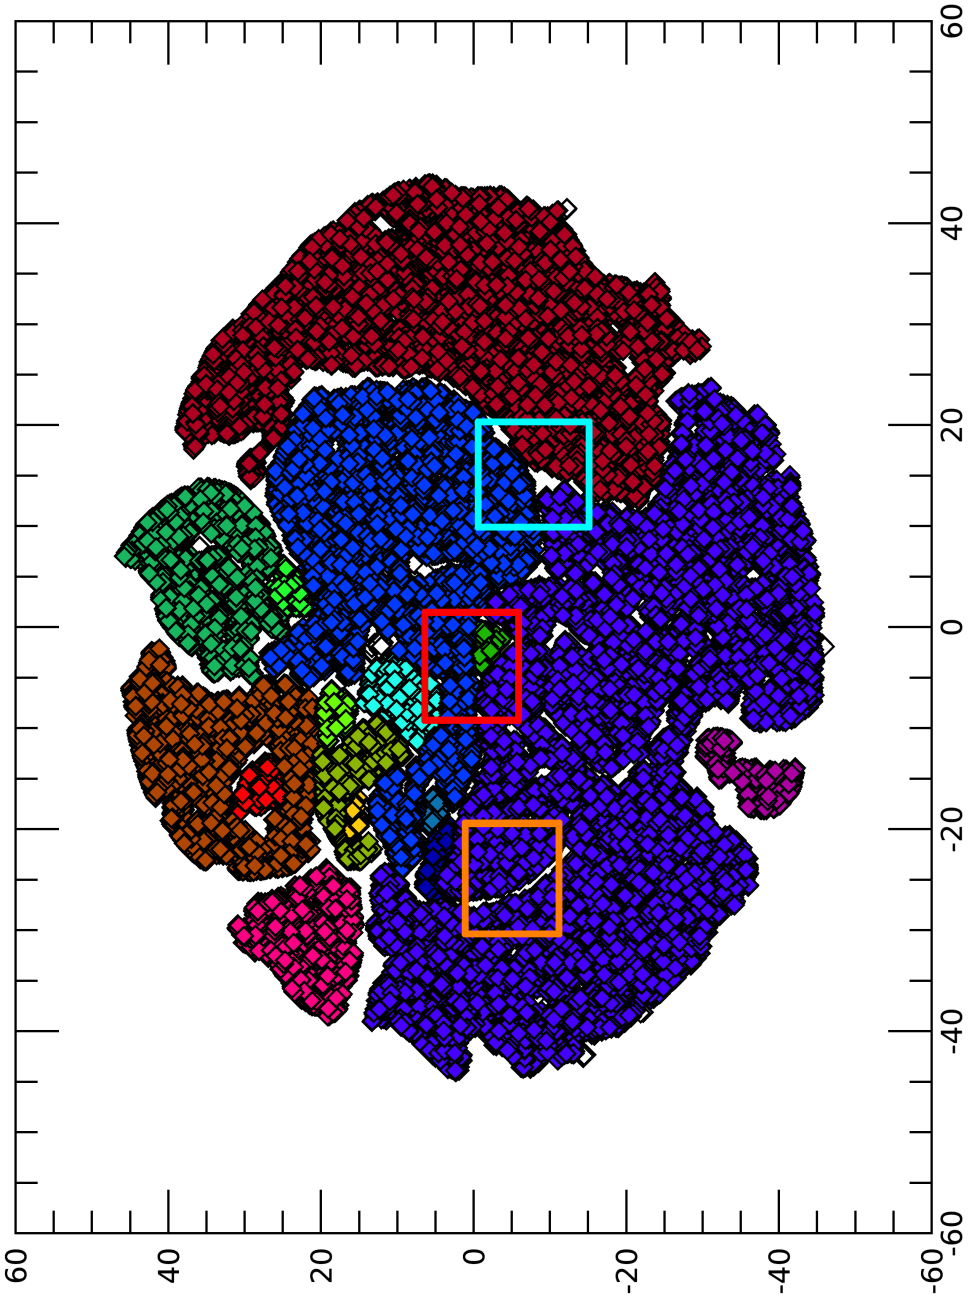

Figure 5.9: t-SNE map. Perplexity of 50.

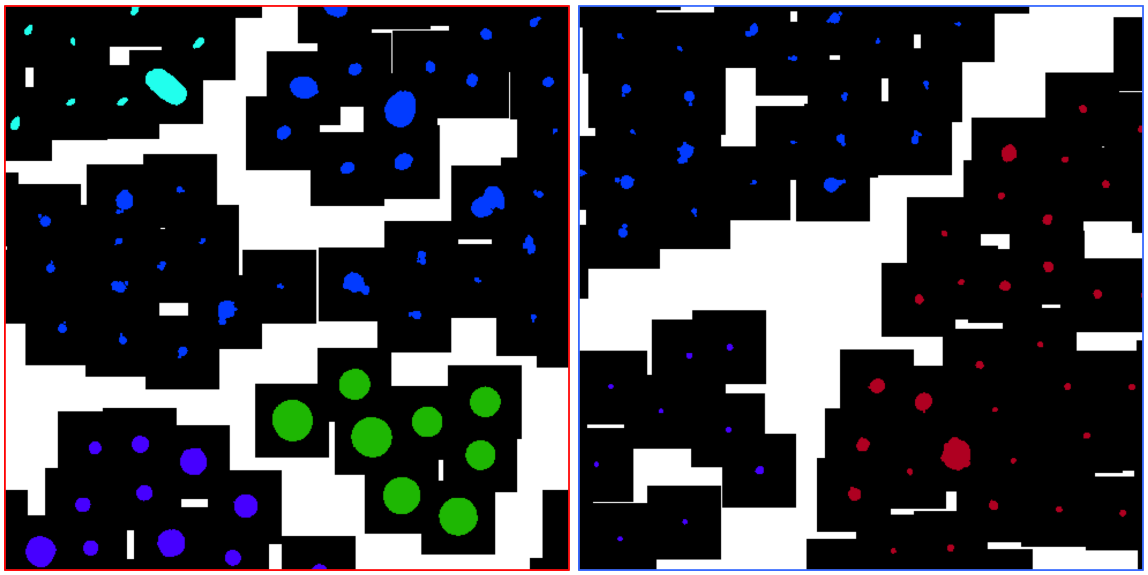

(a) Particle example, red bordered region.

(b) Particle example, blue bordered region.

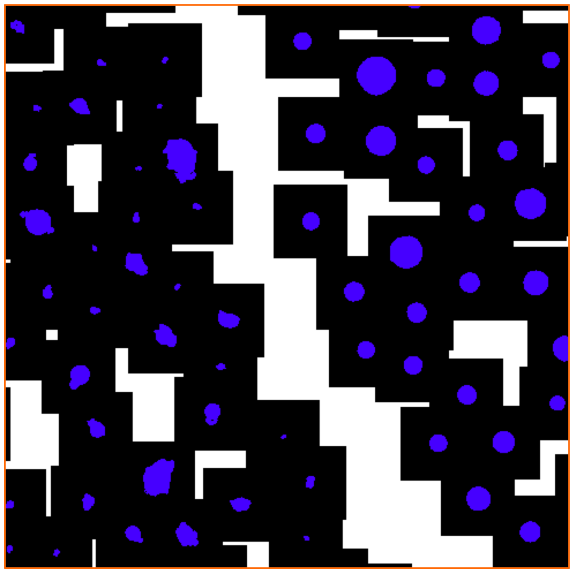

(c) Particle example, orange bordered region.

Figure 5.10: Example of results from t-SNE mapping from five-dimensional ω-space to two-dimensions.

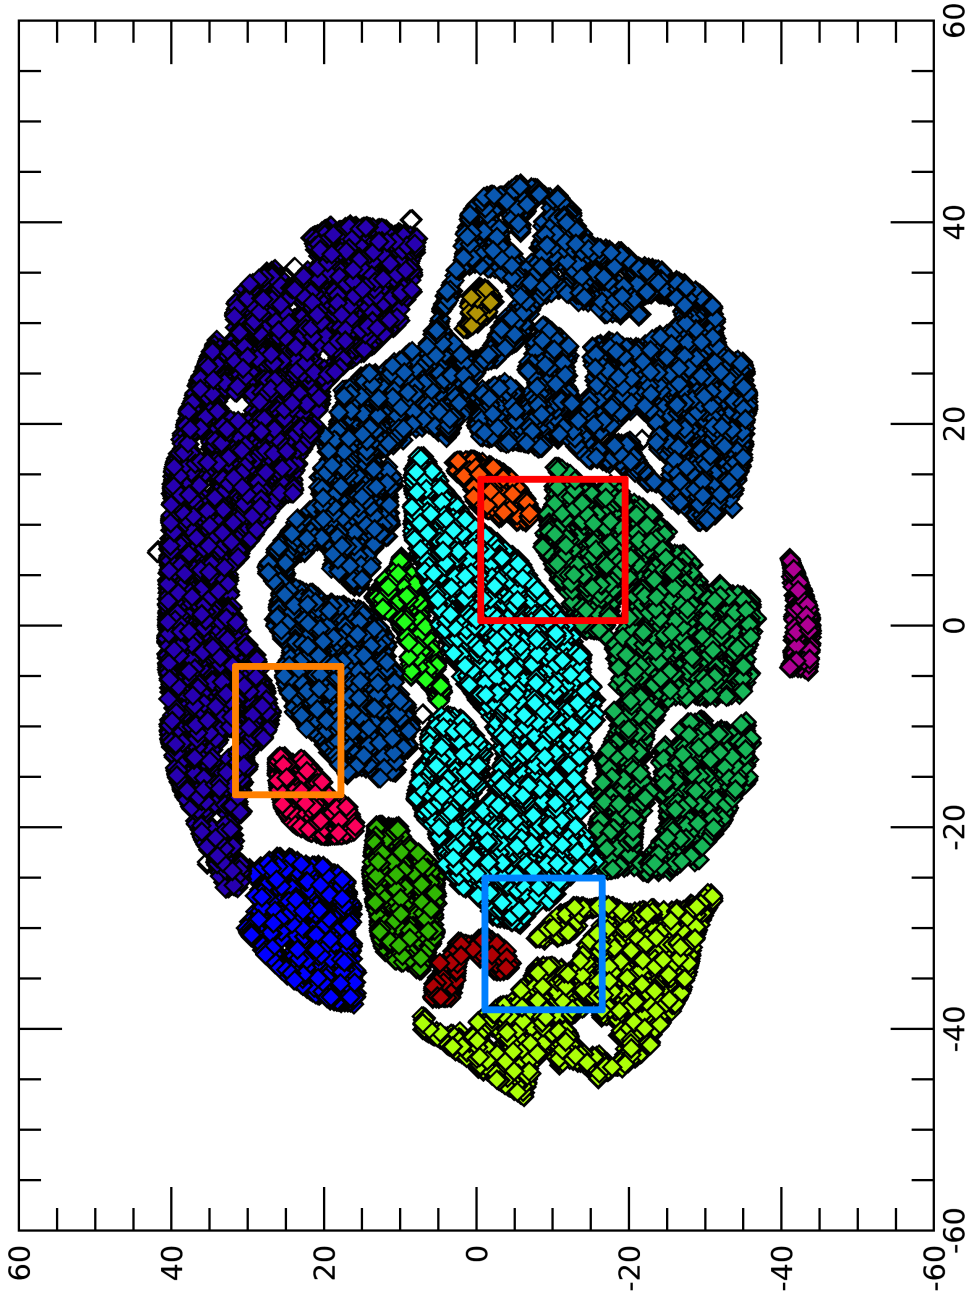

Figure 5.11: t-SNE map. Perplexity of 50.

Confusion matrices for the results of the k-NN generated bag-of-words matching are shown for the ω-based methods in Figures 5.13 through 5.16. It appears that matching to overall bag-of-words histograms for a material is more successful than matching to individual training images, with relative success rates of approximately 86% to 80% respectively. In addition, it appears that distance-based voting for the k-nearest neighbors classifier does not improve the classification. This method appears to have particular difficulty with the classification of maraging steel powder images, with the highest success rate only being 70%.

Similar confusion matrices for the Cartesian invariants are shown in Figures 5.17 through 5.20. These methods showed the best performance of all methods examined here. Once again, matching to entire material histograms appears to be more successful than matching to individual training images. In this case, distance voting had an impact in improving the results of the classification for the matching to bulk materials, but not in the case of matching to individual images. Finally, the maraging steel once again showed that it was the most difficult to classify, commonly being misidentified as In-718.

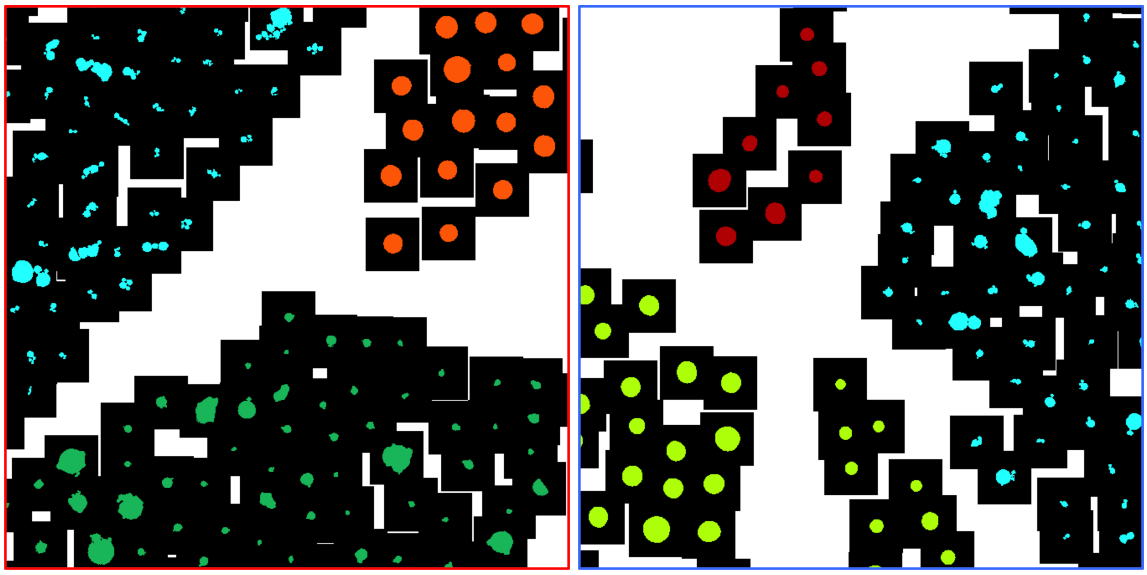

(a) Particle example, red bordered region.

(b) Particle example, blue bordered region.

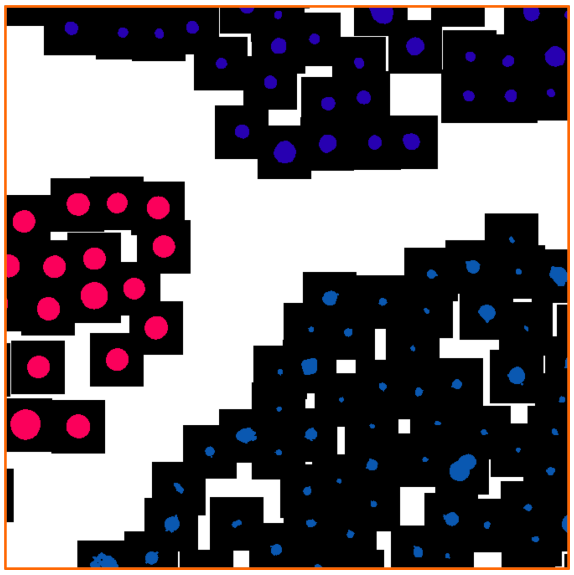

(c) Particle example, orange bordered region.

Figure 5.12: Example of results from t-SNE mapping from eighty-dimensional Cartesian moment invariant-space to two-dimensions.

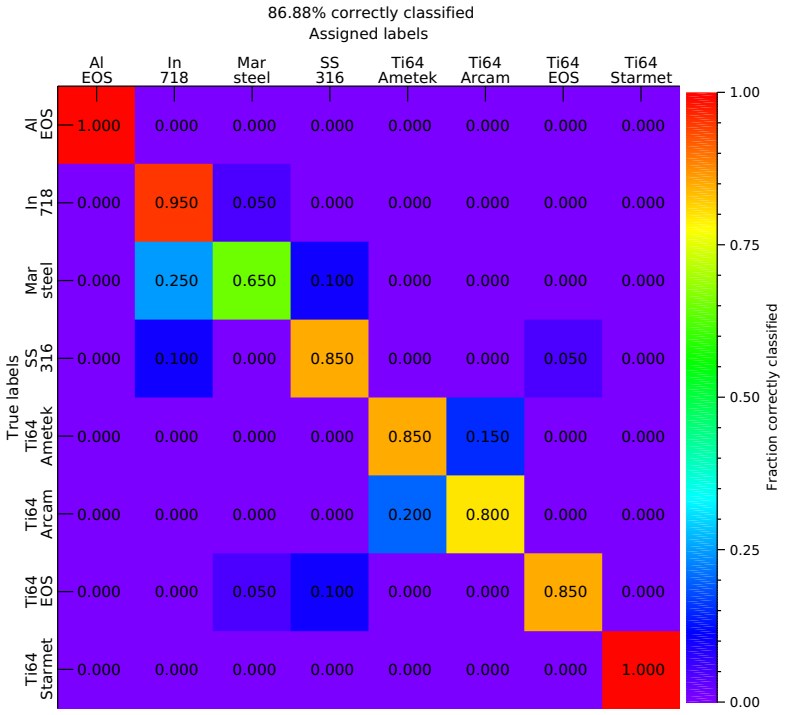

Figure 5.13: Confusion matrix for ω invariants, matching image histograms to material histograms.

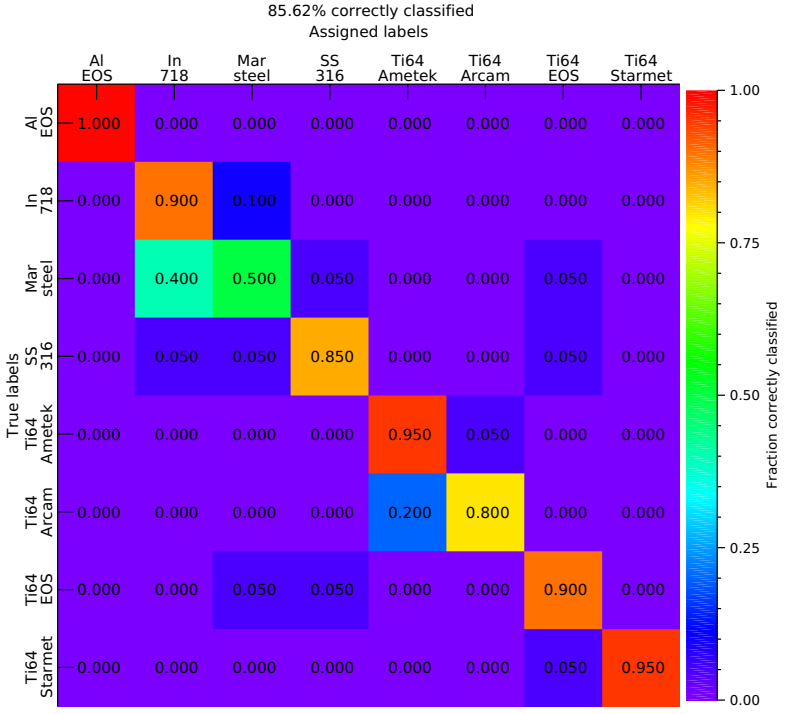

Figure 5.14: Confusion matrix for ω invariants, matching image histograms to material histograms, with distance-based voting for the k-NN classifiers.

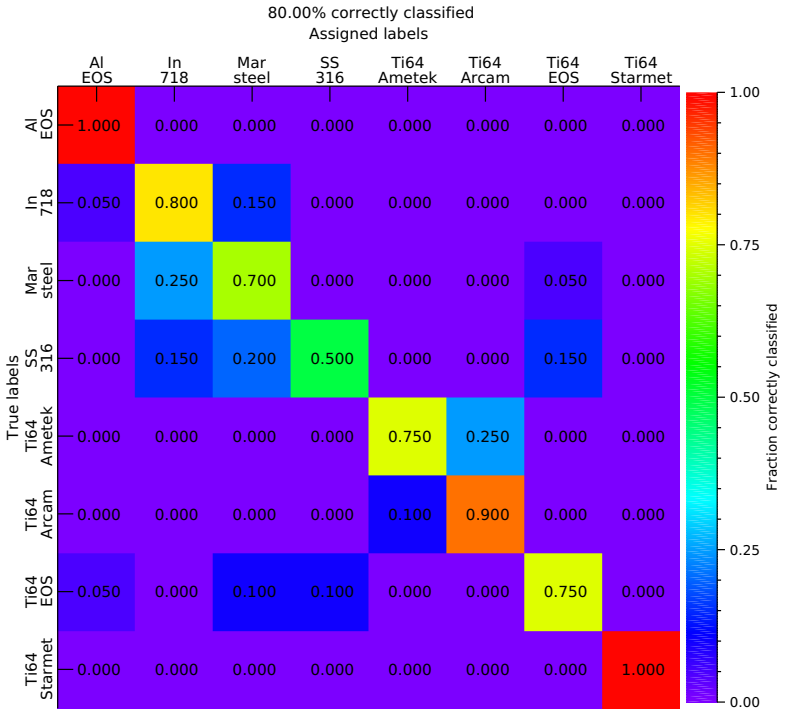

Figure 5.15: Confusion matrix for ω invariants, matching image histograms to training image histograms.

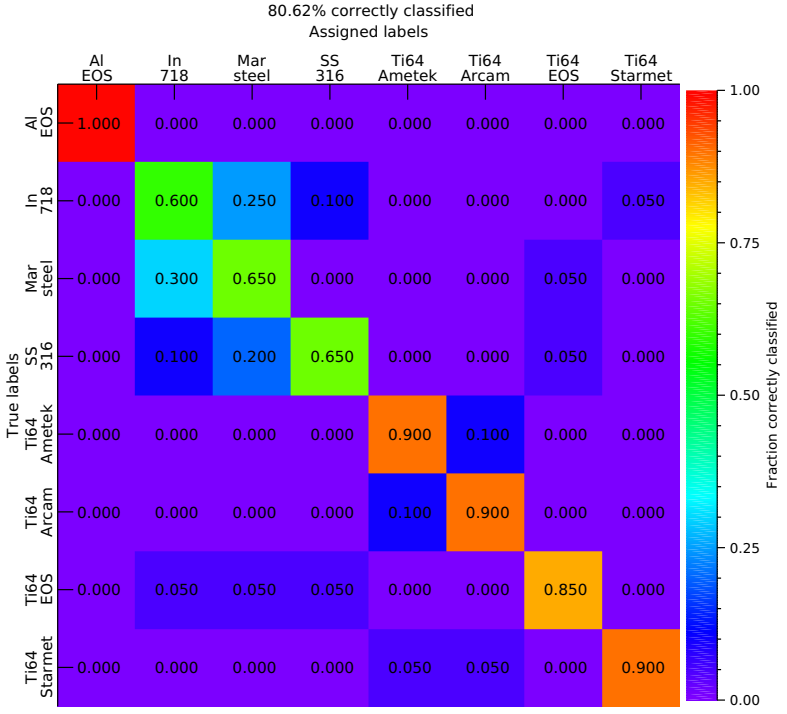

Figure 5.16: Confusion matrix for ω invariants, matching image histograms to training image histograms, with distance-based voting for the k-NN classifier.

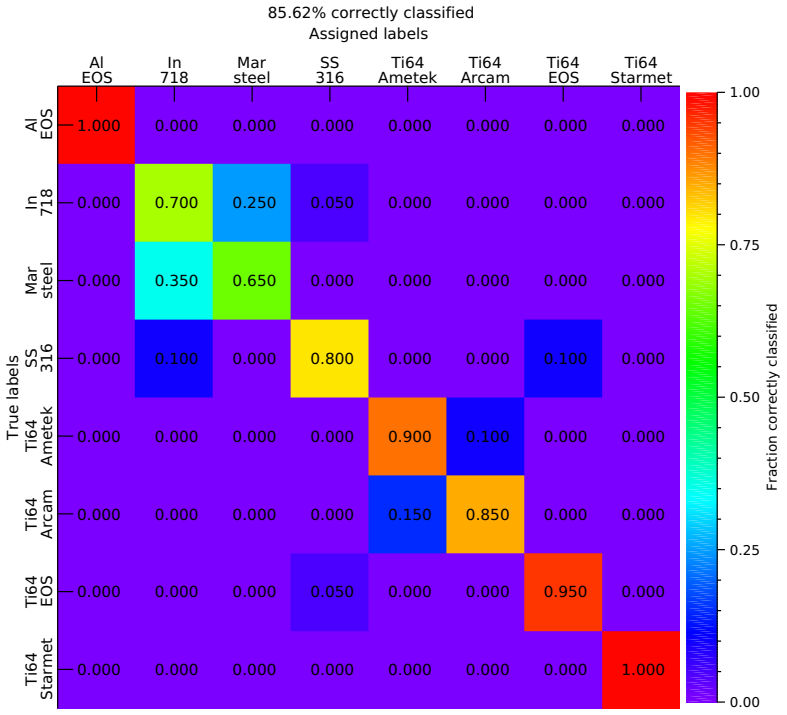

Figure 5.17: Confusion matrix for Cartesian invariants, matching image histograms to material histograms.

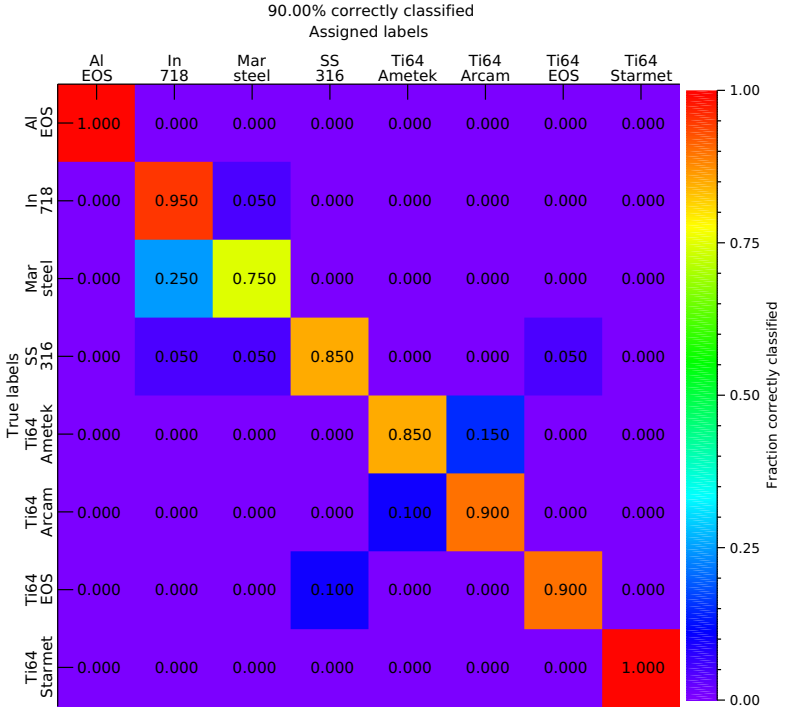

Figure 5.18: Confusion matrix for Cartesian invariants, matching image histograms to material histograms, with distance-based voting for the k-NN classifiers.

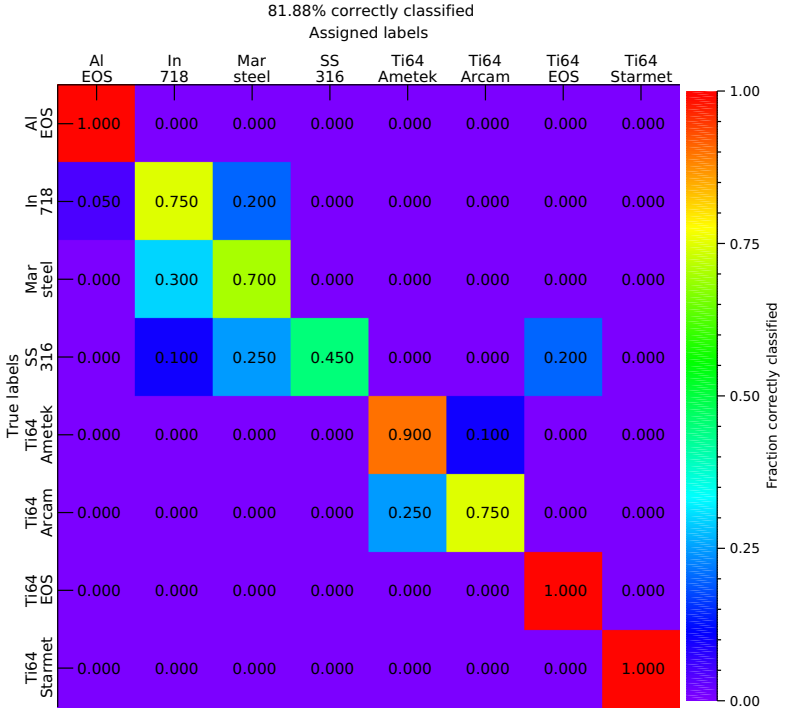

Figure 5.19: Confusion matrix for Cartesian invariants, matching image histograms to training image histograms.

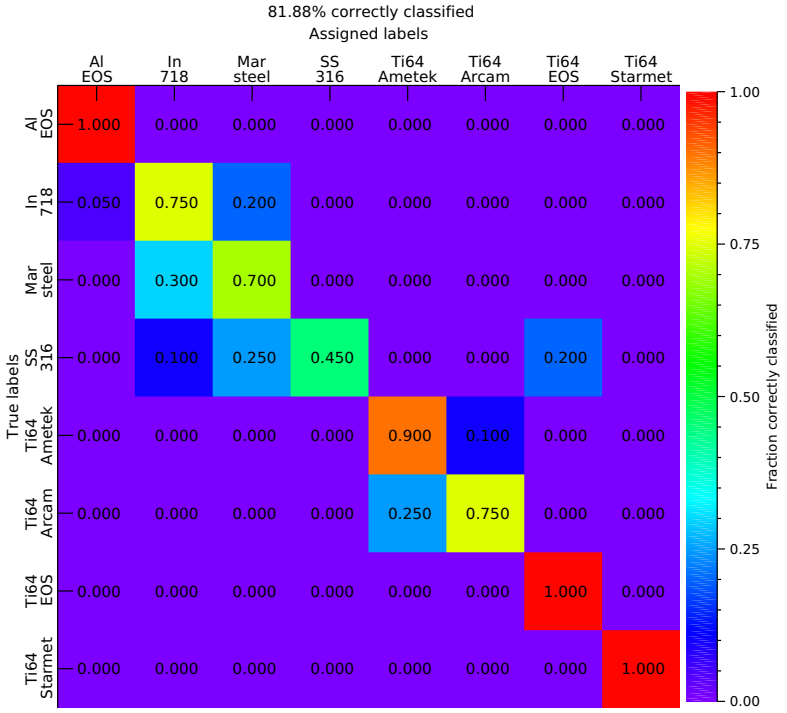

Figure 5.20: Confusion matrix for Cartesian invariants, matching image histograms to training image histograms, with distance-based voting for the k-NN classifier.

5.4 Discussion

5.4.1 Map matching

The success of SOMIM and PMIM matching was clearly only partially successful. The materials which were most successfully classified using the method were Al-EOS, Ti64-Ametek, and Ti64-EOS. Referring to the SOMIM and PMIM for Al-EOS in Figures 5.4a and 5.5a and the micrograph in Figure 5.1a, it is clear that this material has the most extreme variation and morphology between particles. It is likely this extreme spread that makes the map matching method so successful for this material.

Ultimately, the maps, but especially the PMIMs, don't show much difference between each other in two-dimensions. The spreads only seem to differ along a single curve. Because of this, the materials appear to be being classified only according to their relative spread along that single direction. There are three clear delineations: Al-EOS has a large spread; In-718, Mar-steel, SS-316, and Ti64-EOS have intermediate spreads; and the remaining Ti64 powders have relatively tight spreads of shape morphology.

These three levels make themselves obvious in the confusion matrices for this method in Figures 5.6 through 5.8. Al-EOS is classified successfully, due to its uniquely large variance in particle morphology. The intermediate materials show a common degree of misclassification between themselves, with some of the titanium-free materials being classified as Ti64-EOS, another material with an intermediate spread. The three remaining titanium samples are commonly misclassified as one another as well.

The exceptions to these misclassifications are Ti64-Ametek and Ti64-EOS, which both had high success rates. Both of these powders also showed a relatively high degree of having other powder images misattributed to them. Looking at the SOMIMs and PMIMs for these materials makes it somewhat clear as to why. Ti64-Ametek shows the tightest peak with little to no noise: the cross-sections of the particles are all circular or very close to it. The other titanium materials have slight deviations away from this near-perfect peak, which were likely caused by outlier particles, so when the Hellinger distance is calculated between a Ti64 test image with nearly all spherical particles, the master SOMIMs or PMIMs which have deviations will generate a slightly higher distance than those of the Ti64-Ametek

materials. Only for those images which contain outlier particles, which are more rare, that directly match those in the training set, will the distances be lower.

This effect is likely at play in the case of misattributed particles to the Ti64-EOS material. It has a tighter spread near the peak, so test images for In-718, Mar-steel, and SS-316 which have more spherical particles than average can show a closer relation to the slightly better Ti64-EOS than the correct material.

It is clear that the way in which Hellinger distance is calculated directly impacted the results of this test. Hellinger distance is affected by the location of peaks of large numbers of particles. However, by definition, these clusters do not contain outliers, and it is likely in the nature of the outliers that any major distinguishing features of these powders lie. The location of the peaks in these materials is all the same, close to the value of the circle, so there isn't enough distinguishing information in the peaks alone. This is likely why methods which utilize mean moment invariant information, previously attempted and not described here, were ineffective.

5.4.2 Bag-of-words

In contrast to the SOMIM and PMIM matching models, the models which utilized a bag-of-words type classifier were far more successful in the classification task at hand. In general, the classifiers utilizing the more complete descriptors of shape, the affine Cartesian invariants, outperformed those which worked with the lower order ω and τ invariants. Nevertheless, ω invariants were capable of performing with up to 80% success on the 160 test images examined here.

Clustering

The results of the t-SNE mapping from ω/τ space, and HDBSCAN clusterings, visible in Figure 5.9 and, in part, in Figure 5.10, illustrate the effectiveness of t-SNE with the data at hand. Based upon a qualitative examination of the particles separated into their respective clusters, particles are, in general, clustered with those that are visually similar. The teal/light blue region contains particles with elongated major axes. The blue cluster in the center of Figure 5.10a and the top of 5.10b have more extreme morphologies, indicat-

ing either several small satellites, or fewer large ones. The red cluster in the right side of Figure 5.10b are have *slight* deviations from circular. Looking at the affine Cartesian invariant clusterings in Figures 5.12a,5.12b and, 5.12c, the algorithm separates visually distinct particles.

With both feature sets, however, there is a common effect which ultimately is the result of the nature of t-SNE: several clusters contain circular shapes which, upon visual inspection, appear nearly identical. The t-SNE algorithm attempts to balance global structure with local structure. The choice of a relatively high perplexity value of 50 means that the algorithm considers many more neighbors than it would with a low value. This favors global structure, leading to some distortions and moving some similar particles away from each other and into their own clusters. Distance can be misleading in a t-SNE plot, but ultimately, it is proximity, indicating similarity that is most important [185].

In the case of the ω and τ invariant clusterings, circular particles have been misclustered with particles of more irregular morphologies. This is obvious in Figure 5.10c. While the majority of the indigo cluster is comprised of particles with irregular morphologies, there is a region of circular particles. Its unclear what led to a breakdown of the clustering in this particular region, although it may just be due to the way the algorithm has balanced global and local points in its optimization.

Classification

With respect to the ω and τ invariants, looking at figures 5.13 through 5.16, the majority of the materials were successfully classified using the bag-of-words method. The method which involved matching test image bag of word histograms to the overall histograms for the entire material, rather than to individual training images, also seemed to result in more successful classifications.

Figures 5.13 and 5.14 show the result of the bag-of-words method applied to ω and τ invariants, matching to overall materials, and with and without distance-based voting in the assigning of words/clusters to test image particles. Here, it is clear that the maraging steel still creates some difficulty for the method. While the majority of the mar-steel images were correctly identified, between 25 and 40% were misclassified as Inconel 718. There was a

significant increase in the misclassification when the method was switched to distance-based voting, and it was this change that led to a reduction in overall classification percentage, even though other materials, such as Ti64-Ametek saw improvements with the switch. This suggests that the bag-of-words histograms for In-718 and mar-steel are very similar. Some of the test image particles, with the addition of distance voting, are switching clusters, changing the test image histograms just enough to lead to an incorrect classification. This is also likely in the slight reduction of the success in the classification of Ti64-Starmet. This is supported by the overall similarity in the SOMIMs and PMIMs of these materials.

Matching to individual training images actually causes a reduction in the success of the method, with noticeable reductions in the case of the stainless steel alloy and Ti64-Ametek in Figure 5.15. The reduction in the success of SS-316 is primarily due to an increase in misattribution to mar-steel and In-718. Ti64-Ametek sees an increase in misattribution to Ti64-Ametek, but this is ameliorated somewhat by the addition of distance-based voting in test particle cluster identification in Figure 5.16. This reduced success suggests that the histograms of some test images for SS-316 match those of individual training images of mar-steel and In-718 more closely than others of their own materials. The close similarities between the overall SOMIMs and PMIMs for these materials (Figures 5.4 and 5.5), further support that this is the case. Many of the individual images of these materials are so similar that it is only when matching to entire set of images as a whole is there enough information, statistically, to make a correct classification.

The most successful classifications were performed with the affine Cartesian invariants. When matching to bulk material histograms (Figures 5.17 and 5.18), these invariants outperformed the ω and τ invariants of Figures 5.13 and 5.14 only with the inclusion of distance-based k-NN particle classification. It appears that the addition of this distance-based weighting of votes in the algorithm increased the classification percentage of In-718 by 25%, and that of mar-steel by 10%. With the ω and τ invariants, distance-based voting *reduced* the correct classification of these materials. Improvement implies that the particle cluster assignments for test images more closely match those of the overall material with the addition of distance weighting. This suggests that the higher dimensionality of the Cartesian space and the greater details of particle morphology which are captured by these

invariants allow for distance to play a much more important role than in the smaller, less detailed ω and τ space.

Finally, once again, matching to individual images reduces the effectiveness, especially among the group of In-718, mar-steel, and SS-316. These materials have very similar morphological distributions, and it is frequently the case that a match to the bulk material is necessary, simply because the full histogram for the entire material contains more information that allows for a better match. In a way, the matching to individual training images is similar to a k nearest neighbor classifier with a k of 1. However, these classifiers, especially those with a low k value are sensitive to noise [127], and variations in cluster assignment, especially in the case of the circular particles which have been shown to be split among multiple classes, could prove difficult for this type of classifier.

5.4.3 Comparing approaches

It is clear that the histogram matching method was far less successful than the bag-of-words method in this approach, with the BoW model outperforming the histogram method by 25%-30%. The SOMIM and PMIM matching method is limited in the amount of shape information that it quantifies. However, based upon the relative success of the BoW model using this classifier, the reduced success of the map matching isn't due to the reduced levels of shape information. The ω and τ invariants do have enough shape information to make meaningful distinctions between the shapes in question. The breakdown must therefore be in the method.

Ultimately, as previously mentioned, the weakness of the map matching algorithm is responsible for the success of the bag-of-words model. The primary distinguishing features between these powders are the outliers. The average shape of these particles is somewhat close to circular, which is clear from observing the SOMIMs or PMIMs. The big advantage of the BoW model is that the generation of the words puts added emphasis on the outliers. As seen in the particle images of Figures 5.10 and 5.12, the irregular particles, statistical outliers in all of the materials, get broken up into particles which display similar morphological features. These clusters are then directly used in the determination of similarity. It is this that drives the success of the BoW model in this situation.

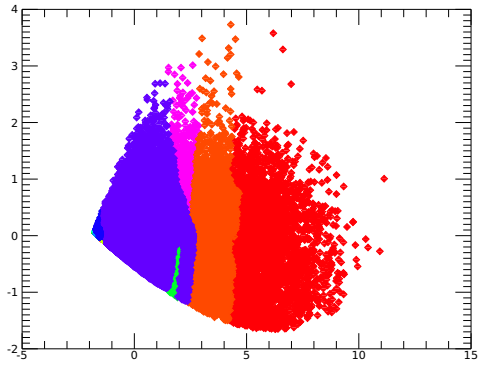

(a) PCA mapping of particles distributed in ω/τ space. Colored by HDBCAN clusters of Figure 5.21c

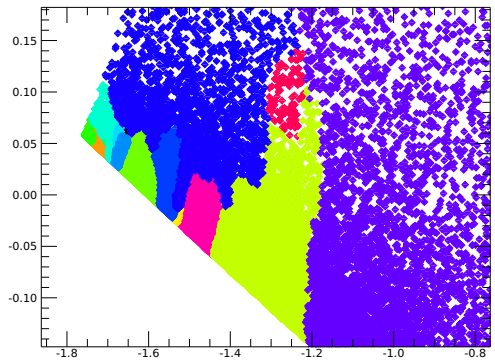

(b) Magnification of the very tip of Figure 5.21a

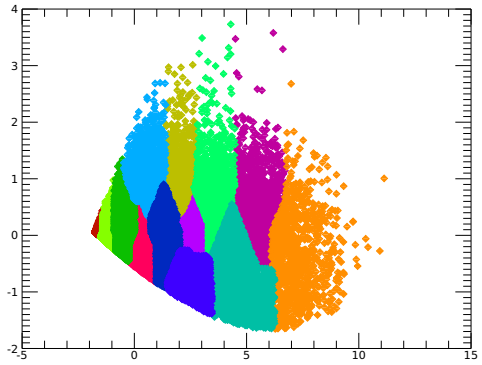

(c) k-means clustering of the same data

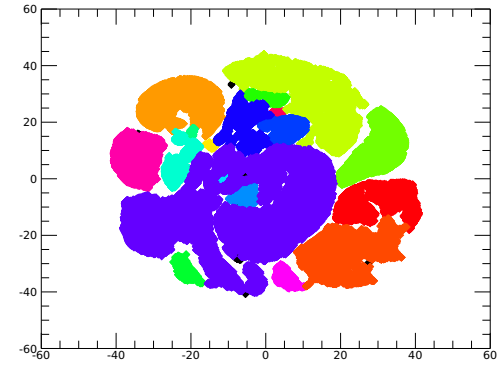

(d) t-SNE mapping and HDBSCAN clustering of particles distributed in ω/τ space.

Figure 5.21: Mapping t-SNE and HDBSCAN clusters back to a PCA mapping, giving an idea of the location of these clusters.

As previously mentioned, it is therefore the ability of the t-SNE algorithm to take an apparently continuous distribution of points, such as the distribution of particle morphologies in the ω/τ space, and from that find underlying structure in the data, that is the source of the success of this method. Figure 5.21a illustrates a PCA-reduced 2D mapping of the distribution of particle morphologies in ω/τ space. The density of points appears to increase monotonically as one moves away from the sharp corner in the bottom left, showed magnified in Figure 5.21b. The particles have been colored according to the t-SNE and HDBSCAN mapping of the data in Figure 5.21d.

As shown in figure 5.21a and 5.21b, the t-SNE algorithm generates many clusters that

have long axes close to perpendicular to the curve that marks the bottom of the data. Looking at Figure 5.21b, the density of points falls off rapidly moving away from this line, indicating that this most likely corresponds to the long curves that define the PMIM plots of these materials. The clusters, therefore, separate the data into increasingly more irregular shapes.

The tip of this plot, which corresponds to the most dense region, where the most circular particles lie, is clustered into many small regions with t-SNE but not in the k-means clustering. This allows subtle differences in circularity to also drive some of the matching. Overall, there may be too many of these small clusters, which shows in the particle maps with many separate clusters of apparent identical circularities, but overall the method is successful.

5.4.4 Comparisons and applications

Compared to the work of DeCost *et al.*, which applied image features to the BoW classifier, this method was less successful. The use of image features allowed for a success rate in excess of 95% on the same data [181]. The success rate of the image feature method against the moment invariant method suggests that there is data that is not being captured by moment invariants. The segmentation of the particles eliminates all of the texture information for the particles, which is stored in the greyscale values across the 2D cross-section of the particle. Image features can use these properties, but moment invariants can only really analyze those surface features which are right at the edge of the particle and change the shape of the border of particle. It is likely that using three-dimensional image with 3D moment invariants would allow for better classification, because all of this information would be captured by the shape features. However, this is unlikely to occur in the near future due to difficulties in gathering 3D information for such a large number of particles.

Generally speaking, the methodology described in this chapter likely succeeded due to the capturing of shape of outlier particles. However, it is clear that it is the most numerically prevalent particles that are going to have the greatest impact on the overall behavior of the powder when it is utilized in an additive manufacturing process. Therefore, despite the fact that this classifier has been shown to be effective, perhaps the lesson of largest importance

to be learned from this work is that the ω and τ invariants contain enough information to make meaningful distinctions between particles of the type that are seen here.

For a technician wanting to manufacture an AM part, they are less likely to be concerned about BoW classifiers or t-SNE or perplexity, and more interested in knowing if a powder is either "good" or "bad". Knowing that the ω and τ invariants are capable of measuring necessary morphological properties, the use of these five quantities, all of which range over $(0, 1]$ (a value of one, in each case, corresponds to a circle), provide a good basis for making morphological measurements of these materials, and are easily interpreted as well. Looking forward, more work must be done on the implications of the value of these quantities on the rheological behavior of particles and their influence on final properties of manufactured parts.

5.5 Conclusions and future work

In this chapter, it has been shown that both the ω and τ invariants, as well as affine Cartesian invariants, are capable of making the kind of distinctions in morphology required to distinguish between similar, and dissimilar, powders for additive manufacturing. In making these distinctions, the most effective shape descriptors were the more complete affine Cartesian invariants up to the 12^{th} order. However, the combination of ω and τ invariants, containing only information to the 4^{th} order, were nearly as effective. A bag-of-words classifier was most effective because it emphasized outliers, primarily driven by the capabilities of t-SNE mapping.

The effectiveness of moment invariants in quantifying powder morphology illustrates their usefulness in this system. Powder morphology has been shown to be linked to the behavior of the particles prior to, and during, the AM process [169–171] and to the packing density of particles, influencing their thermal characteristics [167, 168]. This supports the hypothesis: *shapes of microstructural subregions, as quantified by moment invariants, correlate to the overall properties or function of the material as a whole.*

Looking ahead, more work should be focused on developing a metric, using the two ω values and three τ values, to relate the morphology of AM feedstock to the behavior of the powder during processing and the properties of a final manufactured part. These five values

are easy to understand and have been shown to capture a large degree of necessary shape information.

CHAPTER VI

γ' precipitate shape and creep resistance in single-crystal nickel-base superalloys

6.1 Background

Nickel-base alloys have been the material of choice for high temperature components of jet engines since their creation in the 1940s. However, it was in the 1950s and 60s that the full exploration of the γ' phase, which would become the main source of high temperature performance in modern day superalloys, began in earnest [186]. From the beginning, researchers noticed the distinctive cuboidal shape of γ' precipitates, and in the 1980s, analysis of the evolution of these shapes during aging began for a number of different alloys [187]. While the mechanisms behind the morphology of the precipitates were established at this time, researchers were limited to qualitative observations of particle shapes. This naturally restricted the kinds of conclusions that could be drawn from the data, and prevented the establishment of numerical/quantitative relationships between precipitate morphology and other properties of the superalloy.

6.1.1 Nickel-base single crystal superalloy microstructure and properties

γ matrix

In single crystal nickel-base superalloys, the γ phase forms a continuous matrix. This phase consists of face centered cubic nickel and benefits most from the presence of elements which improve its properties through solid solution strengthening. The elements which are closer in size to nickel tend to concentrate in the γ phase [11]. These are the d-block transition metals. For the samples examined here, the elements are Cr, Co, Hf, Re, W, Mo, and Ru.

These elements have two primary effects on the FCC phase. The first is to improve strength through solid solution strengthening. Local elastic stress fields created by the size discrepancies between the solvent and the solute atoms hinders the motion of dislocations. The strengthening of nickel through the addition of these elements has been shown both experimentally and through atomistic simulations [188, 189]. The presence of an additional strengthening effect with the addition of d block elements has also been noted [188,190,191]. This effect is thought to be rooted in the interactions between the valence electrons of the solvent and solute atoms, and this effect is increased with increasing valence as the periodic table is traversed to the right [188]. These elements also cause a slight change in the lattice parameter of the γ phase, which can contribute to lattice mismatch, or a difference between the lattice parameter of the matrix phase and that of the γ' precipitates.

γ' precipitates

In contrast to the γ phase, the γ' phase is an ordered intermetallic phase which follows the formula Ni_3Al or Ni_3Ti and has the $L1_2$ structure, as shown in Figure 6.1. The $L1_2$ structure is similar to FCC in that there are four atoms per unit cell, however, in this case nickel atoms occupy face centered positions while aluminum atoms occupy the corners. This structure has been confirmed in several different studies through the use of atom probe tomography [192, 193]. The γ' phase is the primary source of the high temperature performance of these alloys, and as such, this phase typically occupies greater than 50% of the volume fraction of the material. It has been shown that the greatest levels of creep resistance occur in single crystal superalloys with γ' volume fractions near 70%, at least with second and third generation alloys [194,195]. While the d-block elements preferentially segregate to the γ phase, these elements can still be present in the γ' phase and substitute according to size. Larger elements such as Ta, Hf, W, and Mo substitute for Al, while smaller elements such as Co will replace Ni [11, 192, 196]. Precipitates of this phase are coherent with the surrounding γ matrix until they reach a maximum size, at which point they lose this coherency. The boundaries between the γ and γ' phase are sharp, on the order of one interplanar spacing [192].

Substitutions on any of the lattice points in either γ or γ' will have an effect on the

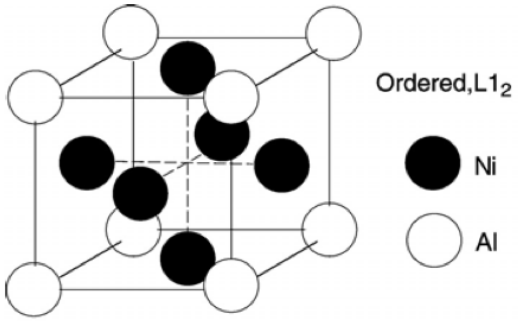

Figure 6.1: The $L1_2$ structure, shown for Ni_3Al. From [11]

lattice parameters of these phases and result in varying levels of misfit (δ):

$$\delta = 2\left(\frac{a_{\gamma'} - a_{\gamma}}{a_{\gamma'} + a_{\gamma}}\right) \tag{6.1}$$

where $a_{\gamma'}$ is the lattice parameter of the γ' precipitate and a_{γ} is that of the surrounding matrix. Changes to the misfit can have a large effect on the nature of the precipitates which form.

The degree of misfit plays an important role in determining γ' morphology. The greater the misfit, the more quickly elastic forces drive the evolution of the precipitate shape [187]. During aging, the γ' precipitates undergo an evolution from spherical to cuboidal to cubic to "octocubic", and finally dendritic, as shown in Figure 6.2. This shape evolution is caused by a balancing of interfacial energy at the surface of the precipitates, and the total elastic strain caused by the misfit between the precipitates and the surrounding matrix. If one assumes misfit between the particles and the matrix remains the same throughout growth, then the total elastic strain is controlled by the volume of the particle. At low volumes, the elastic strain energy is low, so the shape of the particle is controlled by the surface area. The particle forms as a sphere to minimize the surface area to volume ratio.

As the radius of the particle increases, the interfacial energy increases with r^2 while the elastic strain energy increases with r^3. At a particular radius, the elastic strain energy surpasses the interfacial energy and the particle develops a shape that minimizes the energy. The Young's modulus of an FCC material is at its minimum in the $\langle 100 \rangle$ directions, so by growing in these directions, the total strain energy is minimized. With each face growing

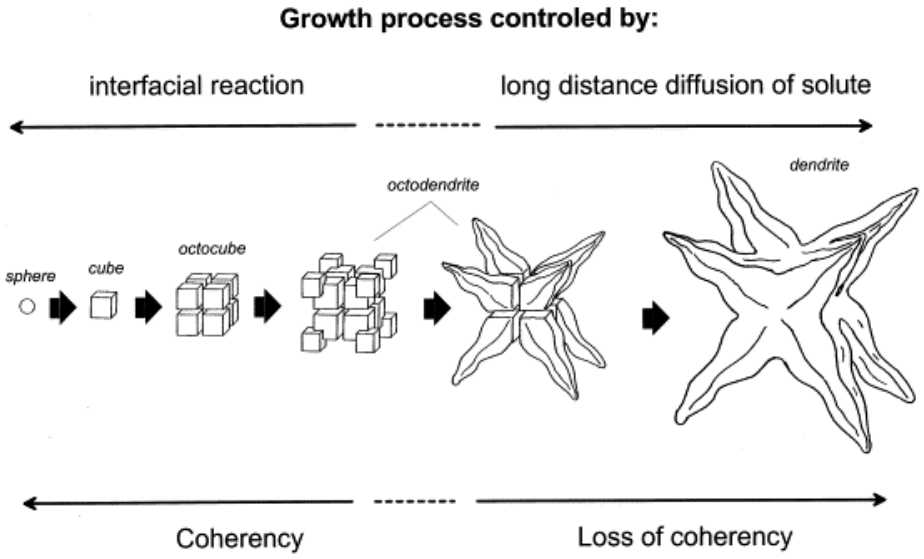

Figure 6.2: The growth of γ' precipitates. From [12]

along these directions, the net effect is growth along the $\langle 111 \rangle$ directions, and the shapes grow at first cuboidal and then fully assume a cubic morphology. Diffusion is accelerated at the sharp corners of the cube, which promotes the growth of spherical lobes. These lobes begin to grow until they are affected by the same elastic restrictions and form cubes. The transition to the more dentritic shapes is thought to occur due to a loss of coherency between γ' and γ [187]. This is intuitive, because the transition from a sphere to a cube is only due to an increase in elastic energy, which can only be generated by the misfit that occurs across a coherent interface. When this coherency is lost, so is elastic energy, and surface energy minimization becomes the driving force behind shape evolution [197].

Mechanical behavior

The two primary properties which engineers and scientists try to maintain at high temperatures in these superalloys are strength and resistance to creep. Strength is simply the maximum stress a material can withstand while not yielding plastically or failing. Creep occurs when a material is placed under a load that is lower than its ultimate strength at elevated temperature. Over time the material will experience deformation in the direction of applied stress. When discussing mechanical performance, it is important to understand the deformation mechanisms in these materials, because the two topics are inextricably connected. Both strength and creep resistance are influenced by the motion of disloca-

tions in the material. By understanding how this motion occurs, it can be understood how superalloys achieve such remarkable high temperature performance.

As previously discussed, the γ phase is face centered cubic, similar to the $L1_2$ structure shown in Figure 6.1, except with nickel atoms at every lattice point. In the FCC crystal structure, the slip system is $\frac{a}{2}\langle 1\bar{1}0\rangle\{111\}$, which indicates that dislocations prefer to move in the $\langle 1\bar{1}0\rangle$ family of directions within the $\{111\}$ family of planes. The magnitude of one of these burgers vectors is equal to $\frac{a}{\sqrt{2}}$, which defines the distance which a dislocation will move in one step, where a is the lattice parameter of the structure in question (in the γ phase this is equal to 0.3517nm [11]). It is well known that the energy attributed to the presence of a dislocation is proportional to the square of the magnitude of the burgers vector. Because of this, in some structures, single dislocations are known to decompose into two dislocations with smaller burgers vectors, known as Shockley partial dislocations. In the FCC structure, this takes place according to the following reaction:

$$\frac{a}{2}\langle 1\bar{1}0\rangle\{111\} = \frac{a}{6}\langle 2\bar{1}\bar{1}\rangle\{111\} + \frac{a}{6}\langle 1\bar{2}1\rangle\{111\} \tag{6.2}$$

When the squares of the magnitudes of the burgers vector above are calculated, the left side comes to $\frac{a^2}{2}$, while the right is equal to $\frac{a^2}{6}$. Therefore it is energetically favorable to assume the two partial dislocations configuration. Also, the magnitude of the burgers vector of each individual partial dislocation is equal to the spacing between close packed planes in the FCC lattice.

The closest packed planes in the FCC structure are the $\{111\}$ planes which stack in the $\langle 111\rangle$ directions. The previously discussed partial dislocations move within the $\{111\}$ planes creating discontinuities in the stacking order, known as stacking faults, by adding planes (an "intrinsic" stacking fault) or removing them (an "extrinsic" stacking fault). Because these defects are a deviation from the tightest possible planar stacking, they have an energy penalty attributed to them. As the two partial dislocation move away from each other, the size of the stacking fault increases, and thus the total energy associated with it increases as well.

It has been shown that the spacing between two partial dislocations is inversely propor-

tional to the energy of the stacking fault between them, as given by [11]:

$$d = \frac{Ga^2}{48\pi\gamma_{SF}}\left(\frac{2+\nu}{1-\nu}\right) \tag{6.3}$$

Where d is the spacing between the dislocations, G is the shear modulus of the material, γ_{SF} is the energy per unit area of the stacking fault, and ν is poisson's ratio for the material. This is because any dislocation, when disassociated, takes on some screw character and some edge character. Depending on the nature of the dislocations in question, they will generate a combination of positive and negative elastic stresses in the lattice, which can attract or repel other dislocations. Equation 6.3 was found by balancing these forces in the disassociated dislocations with the energy of the stacking fault generated between them. In a material that is under a load, moving dislocations create plastic deformation. A screw dislocation which encounters a hindrance to its motion is capable of cross-slipping onto a different slip plane, thereby avoiding the obstacle. If the γ_{SF} is low enough, then the partial dislocations generated will be spread far apart, and thus are unlikely to recombine into a screw dislocation which would be capable of cross-slipping. The alloying elements in superalloys tend to reduce the stacking fault energy and thus work to prevent this method of dislocation motion.

Although the differences between the γ and γ' structure appear to be somewhat insignificant, the presence of aluminum on one fourth of the lattice sites has a significant effect on the nature of the bonds and the deformation of this phase. Unlike in the γ phase, where the order of the atoms on the lattice had no effect on the motion of dislocations, in the γ' phase, maintaining Ni-Al bonds and avoiding Al-Al and Ni-Ni bonds becomes very important. Because of this, if a dislocation passes through a precipitate along a specific direction, it can disorder the lattice in such a way as to add additional energy to the system through what are known as anti-phase boundaries (APBs). When a dislocation enters the γ' phase, the number of Ni-Al bonds decreases, and the number of Ni-Ni and Al-Al bonds increases. This is not energetically favorable. When a perfect dislocation in the γ phase ($\frac{a}{2}\langle 1\bar{1}0\rangle\{111\}$) passes through a γ' precipitate, it leaves an APB in its path, which increases the energy of the system. However, if a dislocation with twice the magnitude ($a\langle 1\bar{1}0\rangle\{111\}$) were to

pass through, the ordering would be maintained and no APB generated. This also implies that regular perfect dislocations can pass through the γ' precipitates if they move through in pairs.

In the γ phase, the primary consideration when determining the nature of dislocations is minimizing the lengths of the burgers vectors in order to minimize the energy of the dislocations. In γ', this is only part of the equation. In addition to burgers vector length, there is also the consideration of the types of previously discussed planar defects which will be generated by the motion of the dislocations. As it turns out, the active slip system in γ' is $a\langle 1\bar{1}0\rangle\{111\}$ [198]. This is twice as large as the system in the γ, and contributes to the complexity of the dislocation decompositions. A number of disassociations have been suggested, and the most common are as follows (on the $\{111\}$ family of planes):

$$a[\bar{1}01] = \frac{a}{6}[\bar{1}\bar{1}2] + \frac{a}{6}[\bar{2}11] + \frac{a}{6}[\bar{1}\bar{1}2] + \frac{a}{6}[\bar{2}11] \tag{6.4a}$$

$$a[\bar{1}01] = \frac{a}{2}[\bar{1}01] + \frac{a}{2}[\bar{1}01] \tag{6.4b}$$

$$a[\bar{1}01] = \frac{a}{3}[\bar{2}11] + \frac{a}{3}[\bar{1}\bar{1}2] \tag{6.4c}$$

Disassociation 6.4a creates two regions known as *complex stacking faults* (CSF), as well as a length of APB. Disassociation 6.4b creates an APB between the resulting dislocations, and disassociation 6.4c creates an *superlattice intrinsic stacking fault* (SISF), between the two partials. According to Yamaguchi *et al.* [199], who performed an energy analysis, 6.4a is most common when APB energy on $\{111\}$ planes is low. When APB energy is high, 6.4b or 6.4c are more common. However, even more complicated disassociations are possible at high temperatures.

It is obvious, based on previous discussion, that moving a dislocation through a γ' precipitate requires much more energy than moving it through the surrounding matrix due to the ordered nature of the γ' lattice. Estimates of the γ' disordering temperature are above the melting point of the material [200]. This contributes to the remarkable high temperature performance of superalloys. However, it has been observed that the yield strength of these alloys increases with temperature. This effect can be at least partially

explained by the trapping of dislocations in what are known as Kear-Wilsdorf locks. These locks occur when at high temperatures disassociated superpartial dislocations slip onto the cubic planes of the γ' phase. On the {001} planes, these partials would trail APBs, and thus tend not to move and are locked in place [201].

The previously discussed mechanisms make it very difficult for dislocations to move in the γ' phase. In general, a number of experiments have shown that during primary and secondary creep, dislocations tend to move in the matrix phase only, and not in the γ' precipitates [202–204]. These dislocations are primarily of screw character (thus allowing cross-slip), and leave edge character dislocations on the interfaces of the γ and γ' phases. Precipitates with negative misfit generate compressive stresses on their interfaces with the matrix, while the opposite is true of those alloys with positive misfit. These stresses are partially relieved by the presence of these dislocations.

In high temperature creep, creep at or above 950°, precipitates also begin to raft during primary creep. This is the preferential growth of precipitates in one direction over another. The degree of rafting is proportional to the misfit of the alloy [205]. In alloys with negative misfit, as most commercial alloys are, the rafts form perpendicular to tensile axis. These rafts close off matrix channels through which dislocation move, increasing creep resistance even more [206]. In alloys with positive misfit, the rafts tend to form parallel to the tensile axis. This is complicated somewhat by the fact that misfit can change with temperature, so it is important to be considering the misfit at the operating temperature, rather than room temperature. Dislocation networks which form at the γ/γ' interfaces stabilize these rafts by relieving stresses and remain in place as a stable configuration [207].

In work by Harada *et al.* [193], the researchers showed that the optimum creep performance in a number of single crystal superalloys increased with increasing volume fraction of γ', but only to a point. At roughly 70%, a peak was reached beyond which creep performance was reduced [193]. This indicates that it is not purely the presence of γ' which influences creep resistance, but the interface between the two phases.

Taken all together, the γ' phase is integral to the performance of these alloys. The morphology of these precipitates is either influenced by factors which also influence mechanical performance, or plays a role in influencing the performance directly. The morphology is

controlled by the amount of misfit, which also determines elastic stresses on the γ/γ' interfaces, and influences dislocation motion. The morphology determines the shape of the γ channels, which allow for tertiary creep in these materials, and controls the geometry of the interfaces through which dislocations interact with the γ' precipitates.

6.1.2 Shape descriptors in superalloy analysis

A number of studies have been performed on the evolution of γ' precipitate morphology in nickel-base superalloys since their invention six decades ago. The first major work illustrating the transition from spherical to cubic to irregular arrays of cubes was qualitative [187]. Soon after, however, more quantitative work was undertaken. These studies utilized shape descriptors which were created by the authors and primarily focused on cuboid cross sections: shapes which were most likely to exist in fully heat treated alloys as prepared for use in turbine engines. Due to the well understood nature of these precipitates, and the single grain structure of the samples, two dimensional analysis is typically all that is necessary for these materials.

Fährmann *et al.* [208] developed and utilized the parameter S_{ster} in their analysis of the influence of the sign and magnitude of misfit strain on γ' precipitates:

$$S_{ster} = \frac{P_i - P_c}{P_s - P_c} \tag{6.5}$$

where P_i is the perimeter of the precipitate in question, P_c is the perimeter of the circle with equivalent area, and P_s is the perimeter of the square with equivalent area. In addition to this parameter, particle aspect ratio was measured as the ratio of the lengths of the primary to secondary axis of the best-fit ellipse. In addition to S_{ster}, the value S_{SAXS} was calculated from small angle X-ray scattering microscopy. This value functions very similarly to S_{ster} in that it measures the particle's shape as its cross section transitions from a circle to a square, however, instead of using visual images, it is calculated from X-ray diffraction patterns. Utilizing S_{ster}, S_{SAXS}, aspect ratio, and measurements of γ/γ' misfit in five Ni-Al-Mo alloys, the authors were able to track morphology changes with aging time. The relationships between aspect ratio and aging time, as well as S_{ster} and aging time were found

to be linear, with the rate of change correlating to the degree of lattice misfit. S_{SAXS} also behaved as expected, with shape evolution behavior being controlled by precipitate/matrix misfit. When plotted against a dimensionless ratio of the elastic energy generated by the precipitate in the matrix and the interfacial energy, S_{ster} increases, indicating an elastically-induced shape change.

When utilized effectively, S_{ster} proves to be a useful quantity, but it has limitations. First, as discussed by Färhmann, both S_{ster} and S_{SAXS} only provide valuable information on the sphere to cube transition when the aspect ratio of the shapes is close to 1. If this value is exceeded, one needs to change the method of calculation of S_{ster} by replacing the perimeters of the equivalent square and circle with those of the equivalent rectangle and ellipse respectively. This somewhat complicates the calculations, but not to a large degree.

When the aspect ratio exceeds 1.5, however, both shape parameters begin to cease giving useful information. This is due partially to the nature of the parameter which becomes more sensitive to perimeter changes as aspect ratio increases, and partially due to the fact that the shapes with large aspect ratios also tend to be more irregular and may have fluctuations in curvature along their perimeters. These can change the value of S_{ster} and make precipitates which seem similar to the eye have different values.

Overall, the parameters presented in Fährmann *et al.* illustrate both the utility and potential pitfalls of using constructed shape parameters. If the user is aware of the limitations of these parameters, they can be useful and allow for the drawing of conclusions regarding morphology. However, if these descriptors are applied to shapes which are outside of their effective range, they not only become useless but risk being misleading.

Prikhodko and Ardell [209] utilized shape descriptors to measure the influence of low applied stress on the evolution of γ' precipitates in nickel-base superalloys. Specifically, the goal was to track shape evolution of individual precipitates to quantify the influence of applied stress on the spherical to cuboidal transition and the coalescence of particles. In order to achieve this, the primary shape descriptor used, Σ was developed:

$$\Sigma = \frac{A_p - A_4}{0.5708(A_4)} \tag{6.6}$$

where A_p is the area of the particle in question, and A_4 is the area of the largest four-sided polygon that can be contained within the particle. The number 0.5708, equal to $\frac{\pi}{2} - 1$, serves to normalize the value of Σ, rendering it equal to 0 for a square and 1 for a circle. Utilizing a descriptor based on area and the area of the largest circumscribed quadrilateral renders Σ invariant to changes in aspect ratio. This is advantageous in the case of this study because it allows for the quick detection of changes in edge curvature in the precipitates.

Prikhodko and Ardell primarily tracked shape evolution in Nickel-base alloys by comparing Σ values and aspect ratios of particles to their average equivalent radius. They found that applied stress tended to drive formation and growth of precipitates with larger aspect ratios and more square cross-sections. They also found that, with increased average radius comes a lower value of Σ, corresponding to more square, rather than circular, shapes, which is supported by what is known about the interaction of precipitate volume and elastic stress.

Morphology and properties

Despite the previously mentioned efforts to quantify precipitate morphology, these studies have remained relatively few and far between, and have primarily focussed on tracking shape change rather than correlating it to mechanical properties. As mentioned in this section, γ' precipitates are the primary source of the strength of these alloys. In the process of creep deformation, propagating dislocations are restricted to the regions between precipitates, and deform to conform to these regions. Because of these, it is suspected that precipitate morphology plays a role in determining the creep resistance in these alloys [210].

Van Sluytman and Pollock [211] tracked precipitate shapes using the value of η which is derived from the shape parameter ω_2 discussed in subsection 3.1.2. Using this shape parameter, it was found that, while a large variety of shapes on the square-to-circle spectrum were observed, a peak in η was observed at a lattice misfit value $|\delta| \approx 0.4\%$. It was suggested that this specific shape which generates the given value of η is optimal for achieve the high-temperature mechanical properties required of these types of alloys.

In her thesis [206], Rowland noted a relationship between precipitate morphology and the creep resistance of ruthenium-containing nickel-base single crystal superalloys. However,

all the observations of γ' shapes were qualitative. Caron and Khan also noted a qualitative relation, between these properties, with more square-shaped precipitates contributing to creep resistance [212].

6.2 Materials and methods

The purpose of this chapter is to explore the relationships between creep resistance in these materials by quantifying the shape of precipitates in several alloys and comparing these measurements to creep properties. In order to do this, aging tests were conducted to see if shape could be quantified accurately in these materials using the planned method of backscattered electron imaging in a scanning electron microscope. This also allowed for the identification of heat treatment times for the final samples that could equalize the size of the precipitates. Ultimately, due to time constraints, only a subsection of the alloys could be fully examined through creep testing.

6.2.1 Sample preparation

Superalloys were provided by Prof. T. Pollock at the University of California, Santa Barbara. The superalloy identifiers and compositions are as shown in Table 6.1 and the details of their casting given in [213] and [214]. Summarizing the casting process, the ME alloys were cast into 15.9mm diameter cylindrical bars by Professor Tresa Pollock's group in an investment cluster mold developed by PCC Airfoils. They were withdrawn from the mold at a constant rate of 20cm/hr. The two UMF alloys were formed as 15.9 mm diameter cylinders in a 50-kW ALD directional solidification furnace. In both cases, traditional helical grain selectors were used in order to achieve single grains in the body of the cylinders.

Alloy	Al	Cr	Co	Hf	Re	Ta	W	Mo	C	Ni
CE-2	5.50	7.00	7.50	0.00	2.80	8.10	4.30	1.50	0.00	63.30
CE-4	5.80	7.00	7.50	0.31	3.10	3.90	4.70	1.50	0.00	66.19
CE-5	5.80	7.00	7.50	0.00	3.00	3.90	4.80	1.50	0.14	66.36
CE-7	5.50	7.00	7.50	0.27	2.80	8.20	4.50	1.50	0.12	62.61
ME-1	6.80	7.00	7.00	0.14	4.60	6.10	2.40	1.50	0.00	63.96
ME-15	7.10	7.00	2.40	0.15	1.50	6.20	2.40	1.50	0.14	66.51

Table 6.1: Alloys used and compositions of each (weight percent).

Alloys				
CE-2, 5, 7	2 hours, 1232°C →	2 hours, 1322°C →	2 hours, 1342°C →	2 hours, 1347°C →
	2 hours, 1352°C →	18 hours, 1357°C		
CE-4	2 hours, 1261°C →	2 hours, 1291°C →	2 hours, 1321°C →	2 hours, 1341°C →
	2 hours, 1351°C →	2 hours, 1361°C →	18 hours, 13666°C	
ME-1	2 hours, 1232°C →	2 hours, 1300°C →	2 hours, 1305°C →	2 hours, 1310°C →
	2 hours, 1315°C →	18 hours, 1321°C		
ME-15	2 hours, 1232°C →	2 hours, 1305°C →	2 hours, 1315°C →	2 hours, 1322°C →
	2 hours, 1329°C →	2 hours, 1336°C →	18 hours, 1340°C	

Table 6.2: Heat treatments for all alloys examined in this study.

To prepare samples, fifteen millimeter long sections were removed from each of the bulk cylinders using an abrasive cutoff wheel. Samples were further divided using a low-speed diamond saw and subsequently encapsulated in quartz backfilled with argon. Samples were then heat treated according to Table 6.2 and then water quenched followed by aging at 1000°C. Heat treatments were conducted in a Lindberg/Blue high-temperature box furnace.

The samples were ground using increasingly fine grit paper, before a final polish with a suspension of 0.5μm alumina, and then five hours on a VibroMet 2 automatic polisher with 0.06μm colloidal silica. Samples were cleaned with soap, followed by ultrasonic cleaning in acetone, and finally by rinsing with ethanol.

Samples were cut in such a way that a $\langle 001 \rangle$ type direction was perpendicular to the sample surface. All samples were observed in a Tescan Mira3 field emission electron microscope at an accelerating voltage of 20kV and a working distance between eight and nine millimeters. While achieving some contrast between the γ and γ' phases is possible using secondary electrons, much better contrast was found using backscattered electrons. Large image sizes were used in order to facilitate analysis of as many γ' precipitates as possible. At least two large images were taken per sample.

6.2.2 Image processing and segmentation

Gathered images are then preprocessed for segmentation by removing all parts of the image which are not of the material (scale bars, etc) and then leveled. Leveling is necessary because of the long scan times, approaching 20 minutes, that were required to obtain the desired large, high quality images. Over the course of such a long scan, fluctuations in

the microscope which affect the intensity across an image can be expected. Leveling is done by treating the intensity levels of each pixel as a z value and fitting a one or two degree polynomial to this surface, effectively returning a plane or a surface which roughly matches the average intensity of the image. This is then subtracted from the original to return an image which has a level image intensity. This dramatically improves the quality of segmentations. In some cases, average image intensity can change from line to line. In this case, each row of pixels is set to have an average value equal to that of the entire image.

In order to effectively gather shape information across a large number of images, quality automated segmentation is required. To accomplish this, the *ilastik* toolkit software was used [15]. *ilastik* takes a machine learning approach to pixel classification starting with a classifier, *decision trees* which are combined into *random forests* [215] which helps prevent overtraining. Initially, the user must manually indicate pixels which are "class 0" or background and pixels which are "class 1" or relevant, although any number of classes may be indicated. This software provides accurate segmentations of the data, as shown in Figure 6.3.

The performance of the *ilastik* algorithm is compared to a previously utilized method of expectation maximization and minimization of posterior marginals (EM/MPM) described in [216]. The results are shown in Figure 6.4. Although the EM/MPM method matches the hand-drawn "truth" up to 91.5%, the location of the most errors are in the matrix regions between particles. The *ilastik* method has a pixel match percentage of 93.2%, a small difference, but it is correct in the right places and matches the truth much better, with less fuzzy edges and properly identified precipitates.

Individual precipitates are identified as connected regions of pixels with a value of 1.Precipitate size can be measured by examining data files which are generated along with each image as it is captured. This data, in the form of area per pixel in square meters, allows for size calculation by simply adding up the number of pixels per region. Moment values are calculated using the algorithm suggested by Novotni and Klein [10]. These moment values are then used to calculate all invariants necessary from each shape. SOMIMs are generated by placing a small Gaussian peak at the location (ω_1, ω_2) for a given shape. These peaks are added up, and then the image normalized, so the brightest region corresponds to

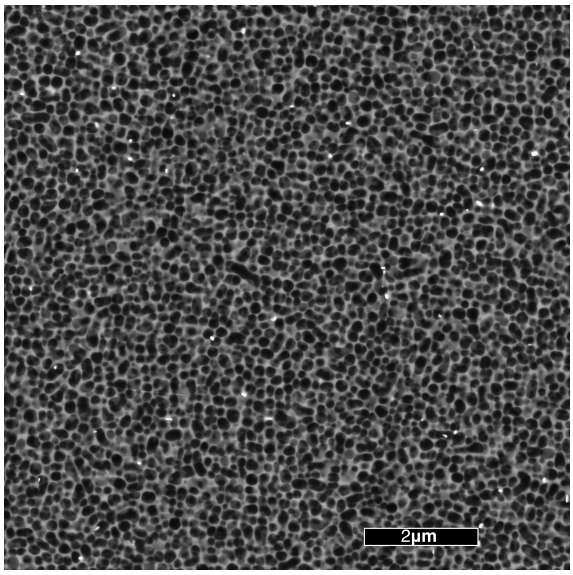

(a) CE-4 alloy at 2 hours of aging, prior to segmentation.

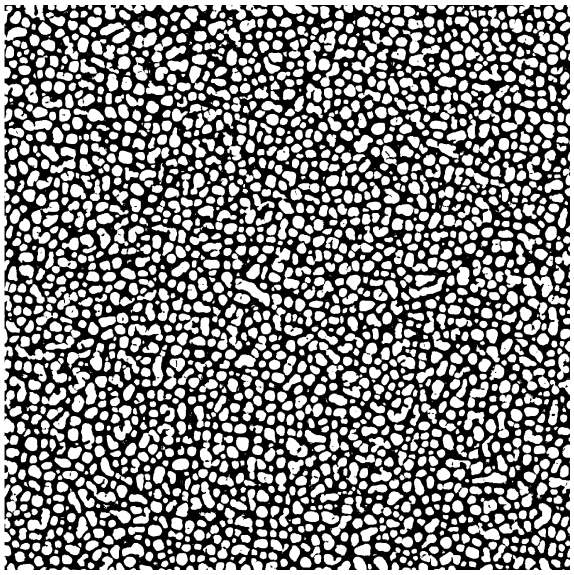

(b) Image from Figure 6.3a, after segmentation.

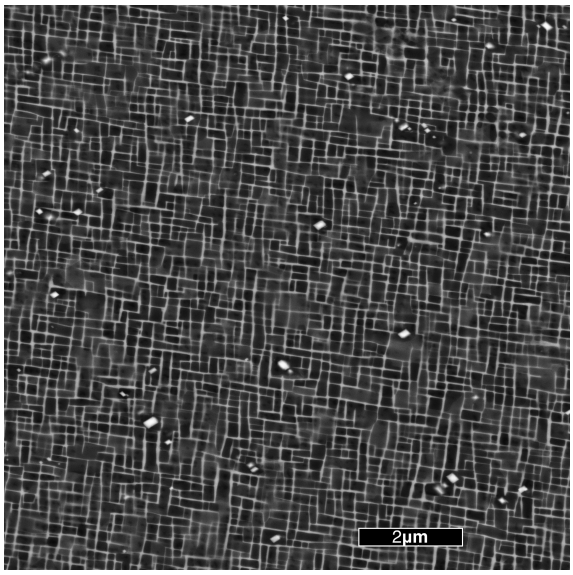

(c) ME-15 alloy at 16 hours of aging, prior to segmentation.

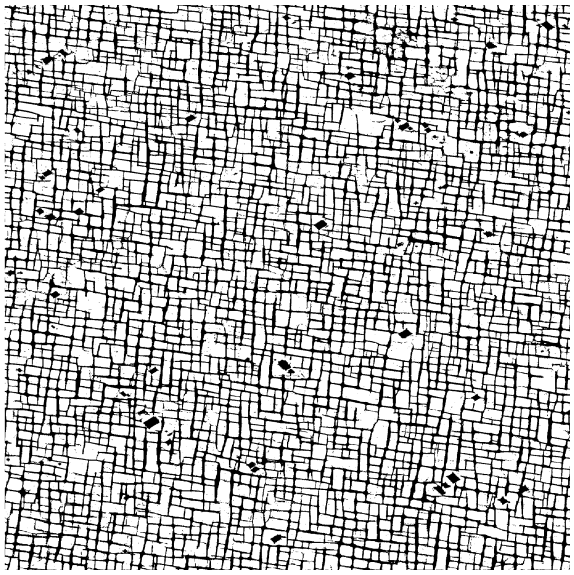

(d) Image from Figure 6.3c, after segmentation.

Figure 6.3: Segmentation examples using *ilastik* software.

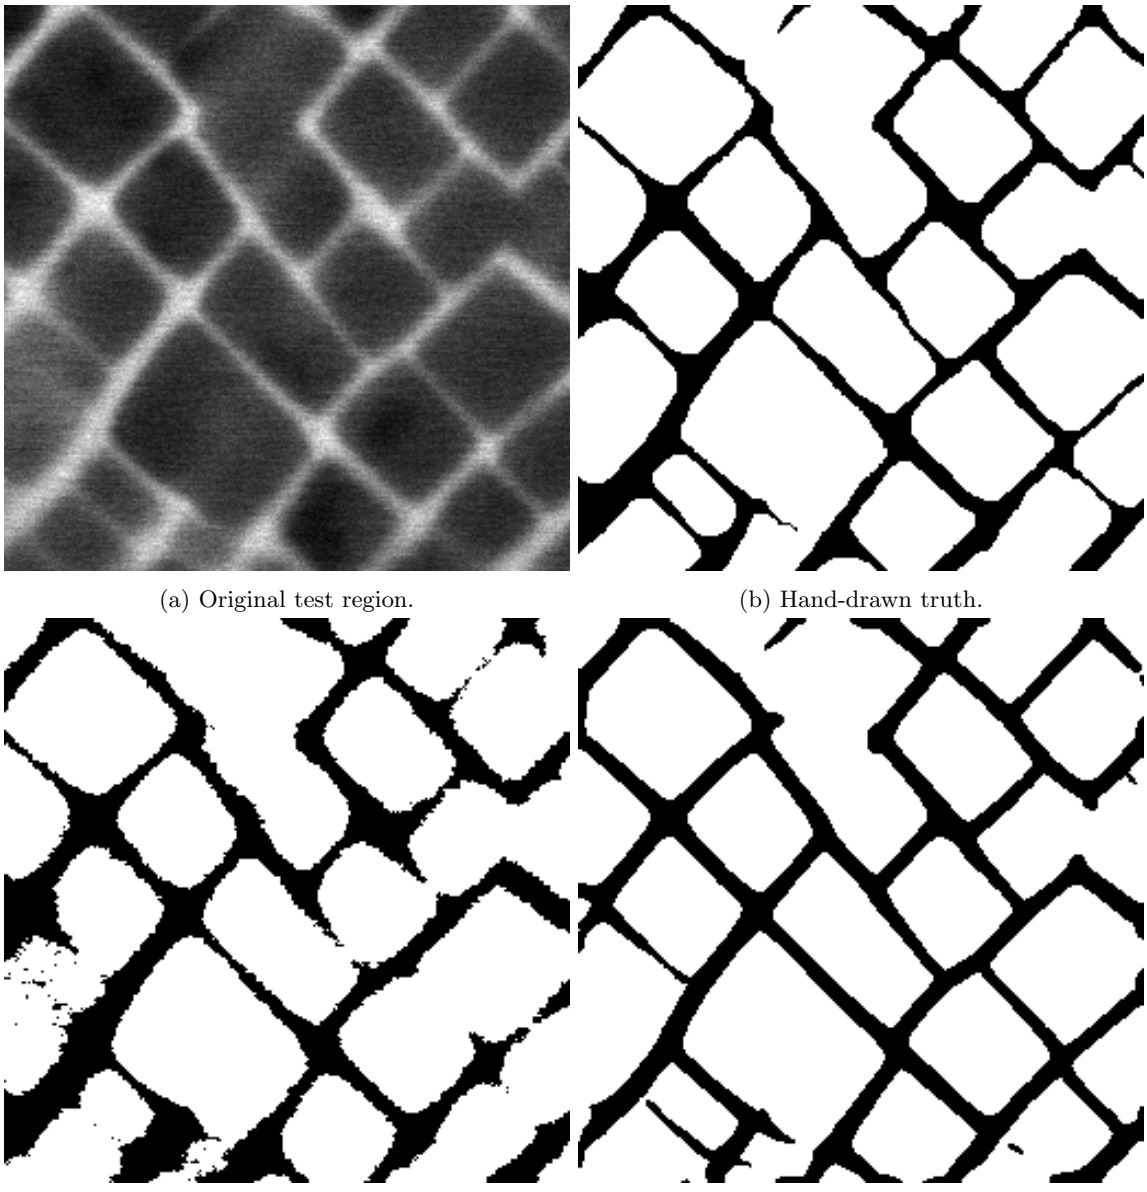

(a) Original test region.

(b) Hand-drawn truth.

(c) EM/MPM segmentation.

(d) ilastik segmentation.

Figure 6.4: Segmentation comparison of EM/MPM algorithm and *ilastik* segmentation method.

locations with more shapes in that local area.

6.2.3 Creep sample preparation

Due to the length of time necessary for creep testing, only four alloys were selected for this analysis. CE4 and CE7 were not subjected to creep tests because of their similarity of composition and final shape to CE2 and CE5. Based upon results from the aging study, described above and results shown below in Figure 6.9, heat treatments were chosen for these alloys. Alloys were solution treated, for 18 hours at 1250°C, and then aged at 1000°C for the following times: CE2: 12.25 hours, CE5: 8.5 hours, ME1: 8.5 hours, ME15: 32 hours.

Due to error on the part of the investigator, samples intended for creep tests were air-cooled rather than water quenched. This lead to precipitate sizes and morphologies that differed from expected values. The implications of this are discussed below, however, there was only a significant change in precipitate morphology for one of the alloys.

These samples were then machined by Metcut Research Inc. (Cincinnati, OH) to a standard creep test specimen with a gauge length of one inch and diameter of 0.145 inches. These were then tested with a stress of 290MPa, at 960°C by Dr. Patrick Callahan of Professor Tresa Pollock's group at The University of California, Santa Barbara.

6.3 Results

6.3.1 Heat treatments and precipitate growth

Images

Figure 6.5 shows the evolution of the microstructure of each alloy with increasing aging time. In the interest of preserving space, only segmented images are shown. Each picture was arbitrarily chosen as a representative from several images taken for each sample and for each heat treatment. Every image of Figure 6.5 is normalized so that the width of the image is the same at $5.18\mu m$.

Visual examination of these images indicates the most obvious influence of aging time on an alloy is precipitate growth and, in the heat treatments after sixteen hours, precipitate coalescence. This is clear from the apparent increase in irregularly shaped particles,

especially those with larger aspect ratios. Aside from the increase in the prevalence of less regular shapes, some general change in shape is also observable, as the shapes appear to become more square with increased aging time. However, based on visual examination, the degree of the transformation is difficult to measure.

Across samples, it is clear that the ME alloys are much more square than the CE alloys, which tend to display more circular or elliptical shapes. Based upon what is known about the influence of elastic and surface energies upon these alloys, this is likely due to the degree of misfit present in these alloys [197].

Precipitate evolution

Second order moment invariant maps (SOMIMs), detailed in Section 3.1.2 and illustrated in Figure 3.4, help to visualize the evolution of the shapes across each heat treatment, and the subtle differences between the shapes present in each alloy. SOMIMs for each alloy and heat treatment are shown in Figure 6.6. It is clear that what is observed visually in the segmented images of Figure 6.5 is supported by the shape analysis. That is that the precipitates of CE type alloys progress toward elliptical or circular shapes, and those of the ME alloys tend to have more square-like cross-sections.

There appears to be a tendency for the peaks on these SOMIMs to become smaller with increased aging times, which indicates an overall increase in the similarity of the shapes in terms of both aspect ratio, and complexity. This does not universally hold, however. In the case of both CE-2 and ME-15, there appears to be a small increase in the spread of the peak. In many cases, the peaks are smeared to the left at higher aging times, indicating that the shapes, while similar in terms of type location along the cuboidal spectrum (square to circle), tend to differ in aspect ratio. Overall, these SOMIMs provide a way to get a snapshot of the shape distribution of a large number of precipitates in a small, easy to understand visualization.

Figure 6.7 and 6.8 present box plots for the data. The shape of the boxes strongly indicate a distributions of shapes with long tails toward irregular shapes (lower ω values). It is probable that these shapes were caused by errors in segmentation. In order to correct for these extreme distributions, trends are drawn between median values, rather than mean

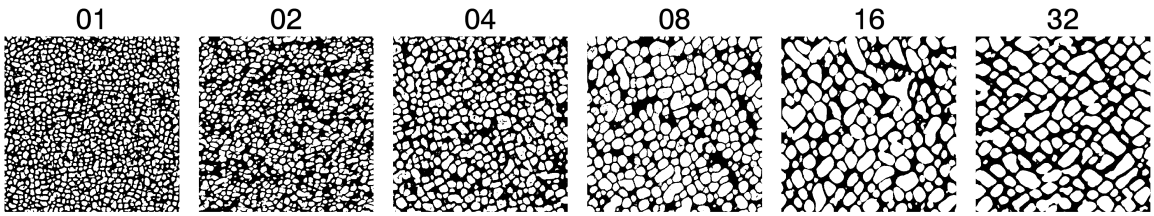

(a) CE-2

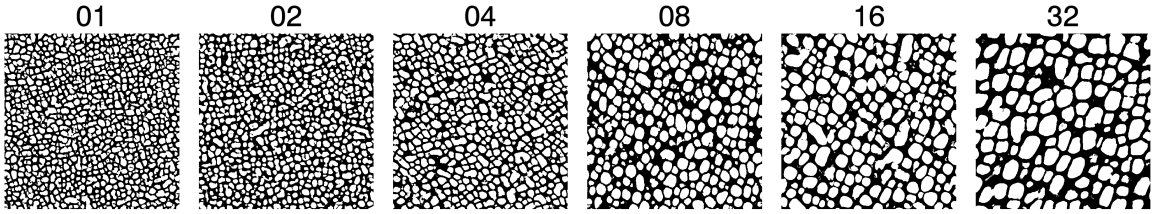

(b) CE-4

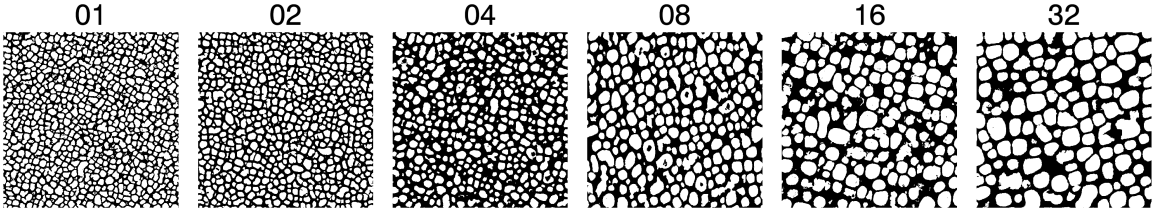

(c) CE-5

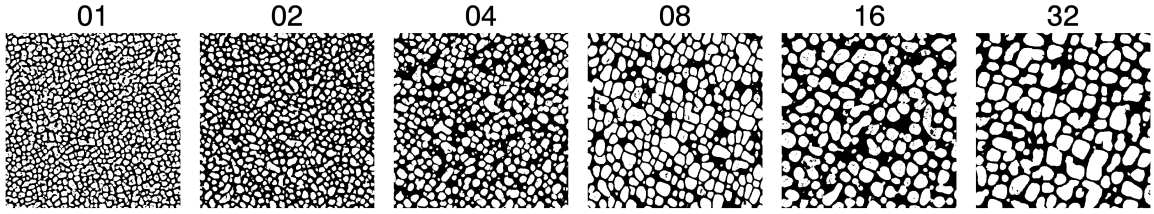

(d) CE-7

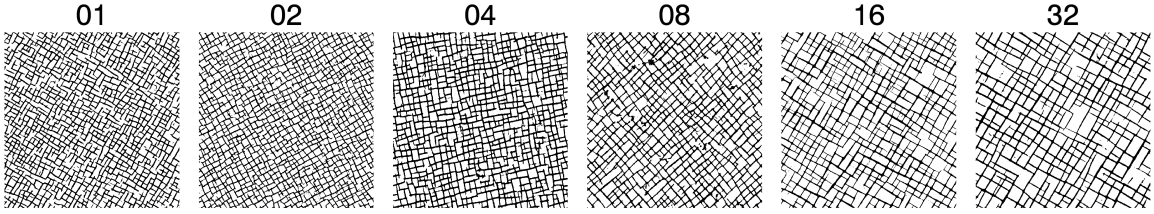

(e) ME-1

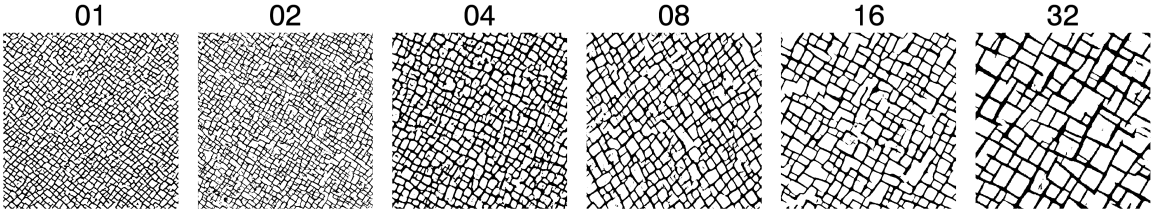

(f) ME-15

Figure 6.5: Segmented microstructural evolution of γ' precipitates for indicated aging times, in hours. All images are $5.18\mu m$ across.

values. The median tends to be more robust against outliers. In general, the materials appear to start with more irregular values and over the course of the first one to four hours approached a fixed shape, indicated by a horizontal line with larger aging times.

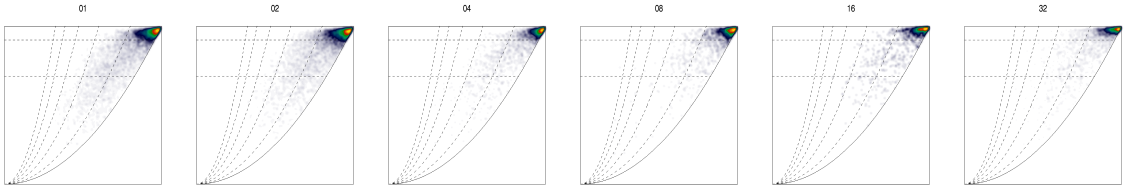

(a) CE-2

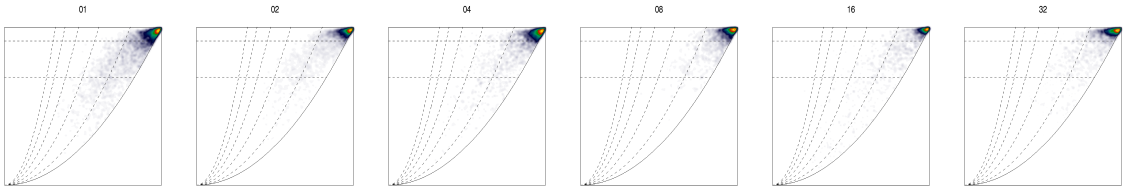

(b) CE-4

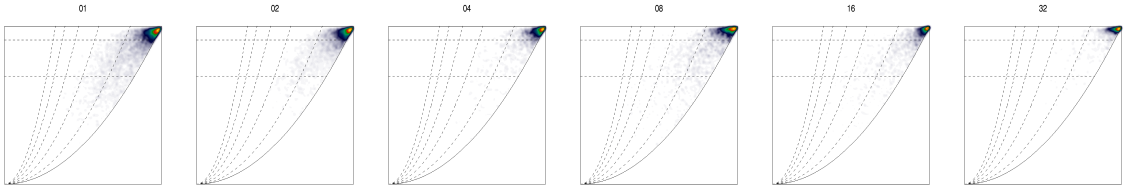

(c) CE-5

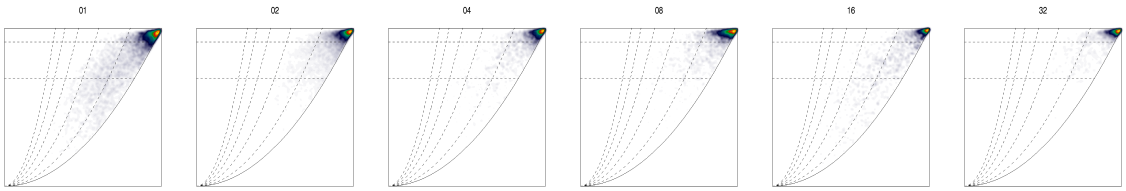

(d) CE-7

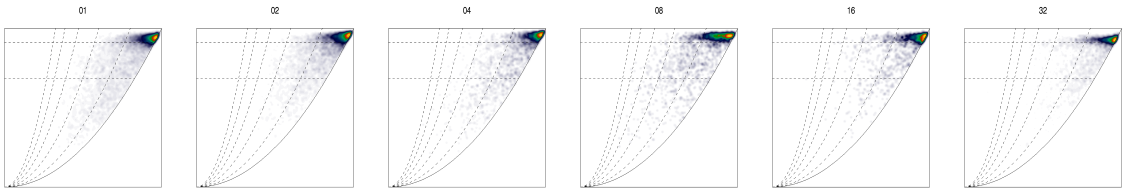

(e) ME-1

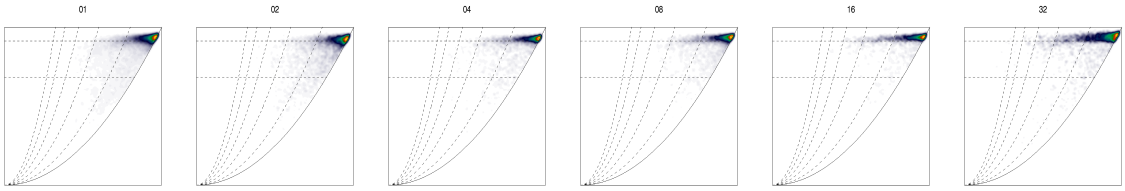

(f) ME-15

Figure 6.6: Second order moment invariant maps for indicated aging times, in hours.

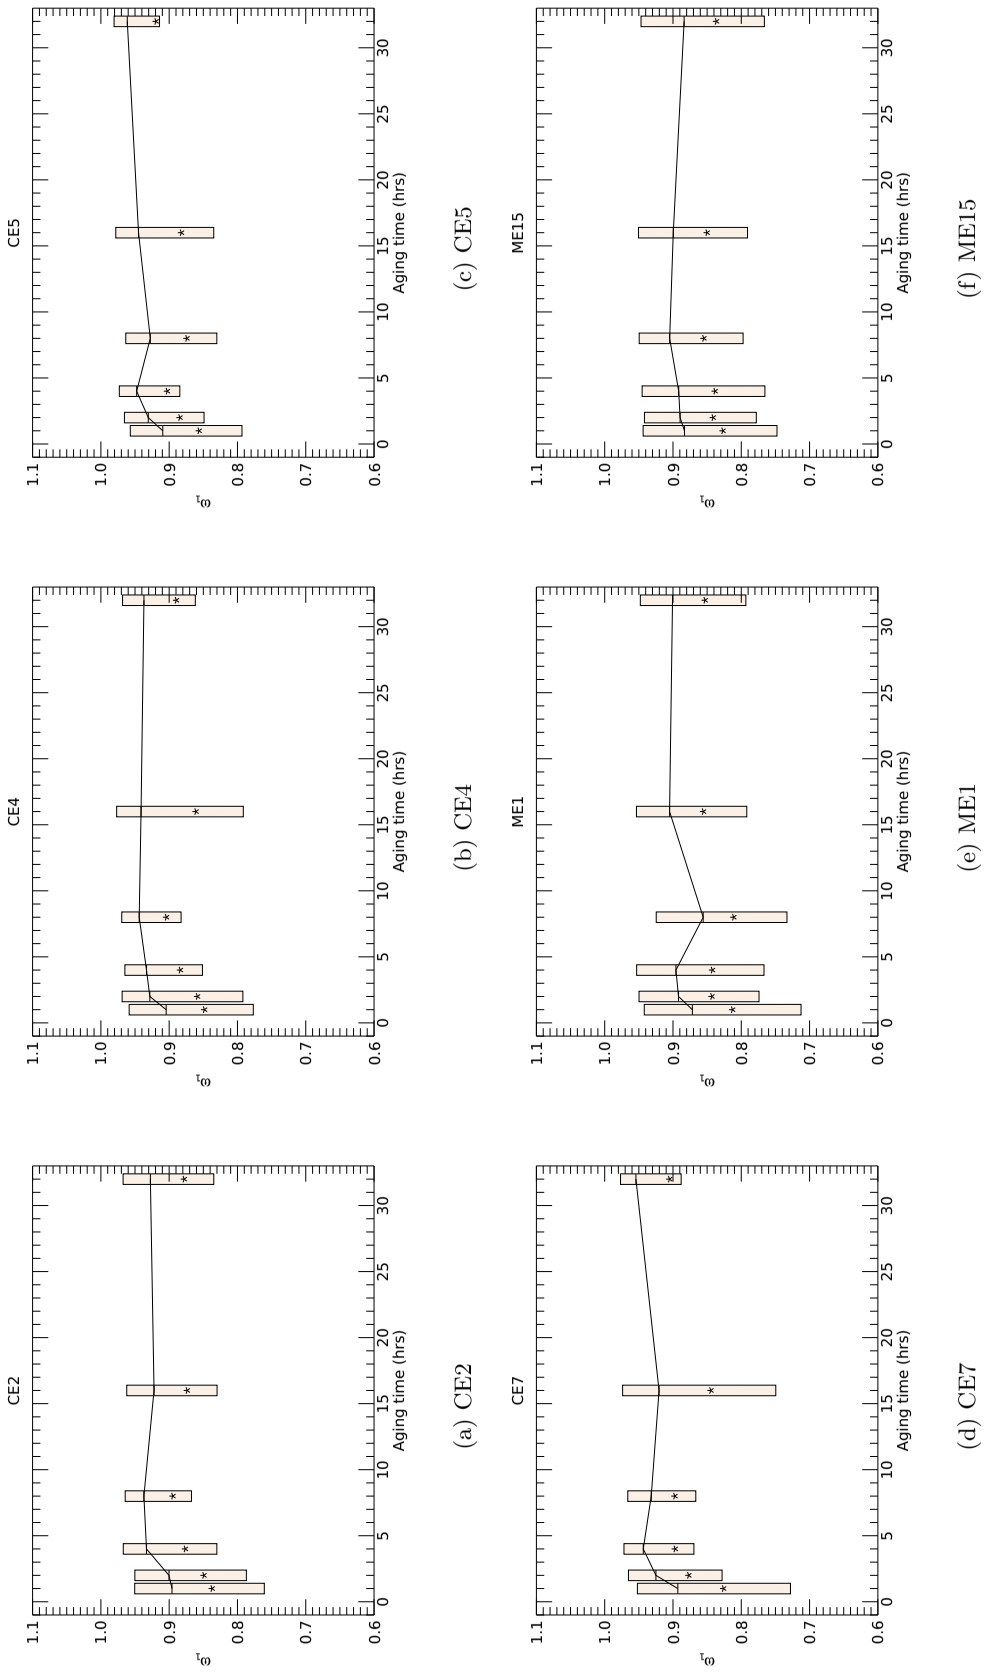

(a) CE2 (b) CE4 (c) CE5

(d) CE7 (e) ME1 (f) ME15

Figure 6.7: Box plots of ω_1 for all measured precipitates. Asterisks indicate mean.

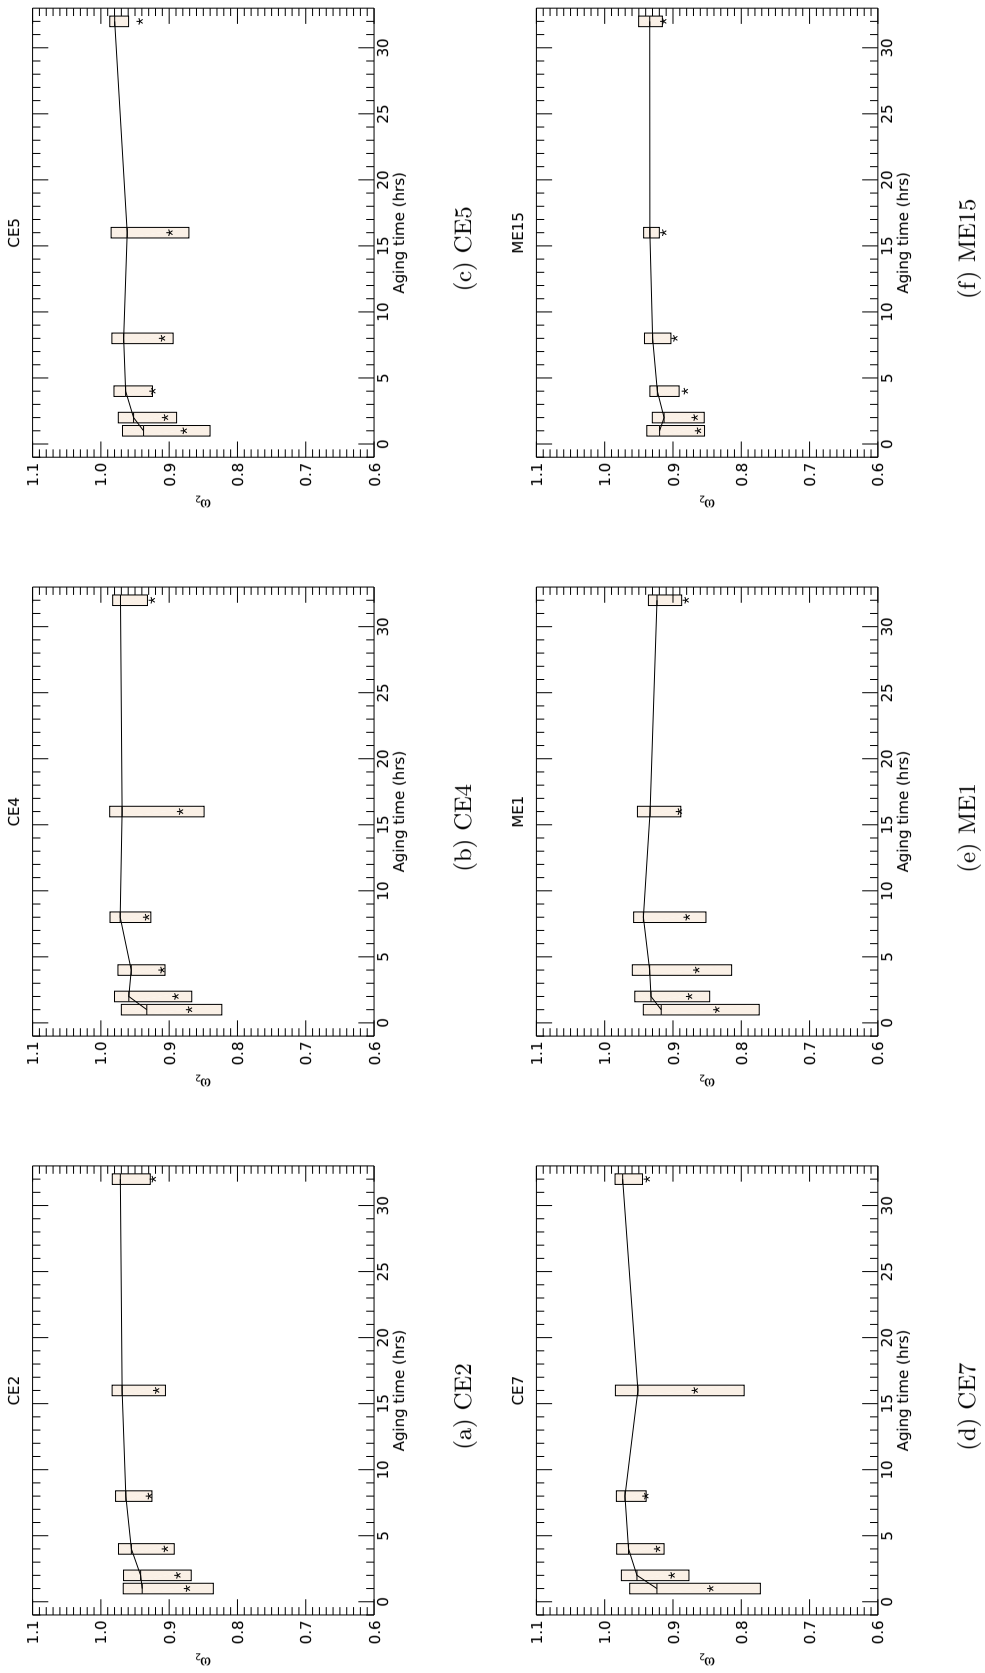

(a) CE2 (b) CE4 (c) CE5

(d) CE7 (e) ME1 (f) ME15

Figure 6.8: Box plots of ω_2 for all measured precipitates. Asterisks indicate mean

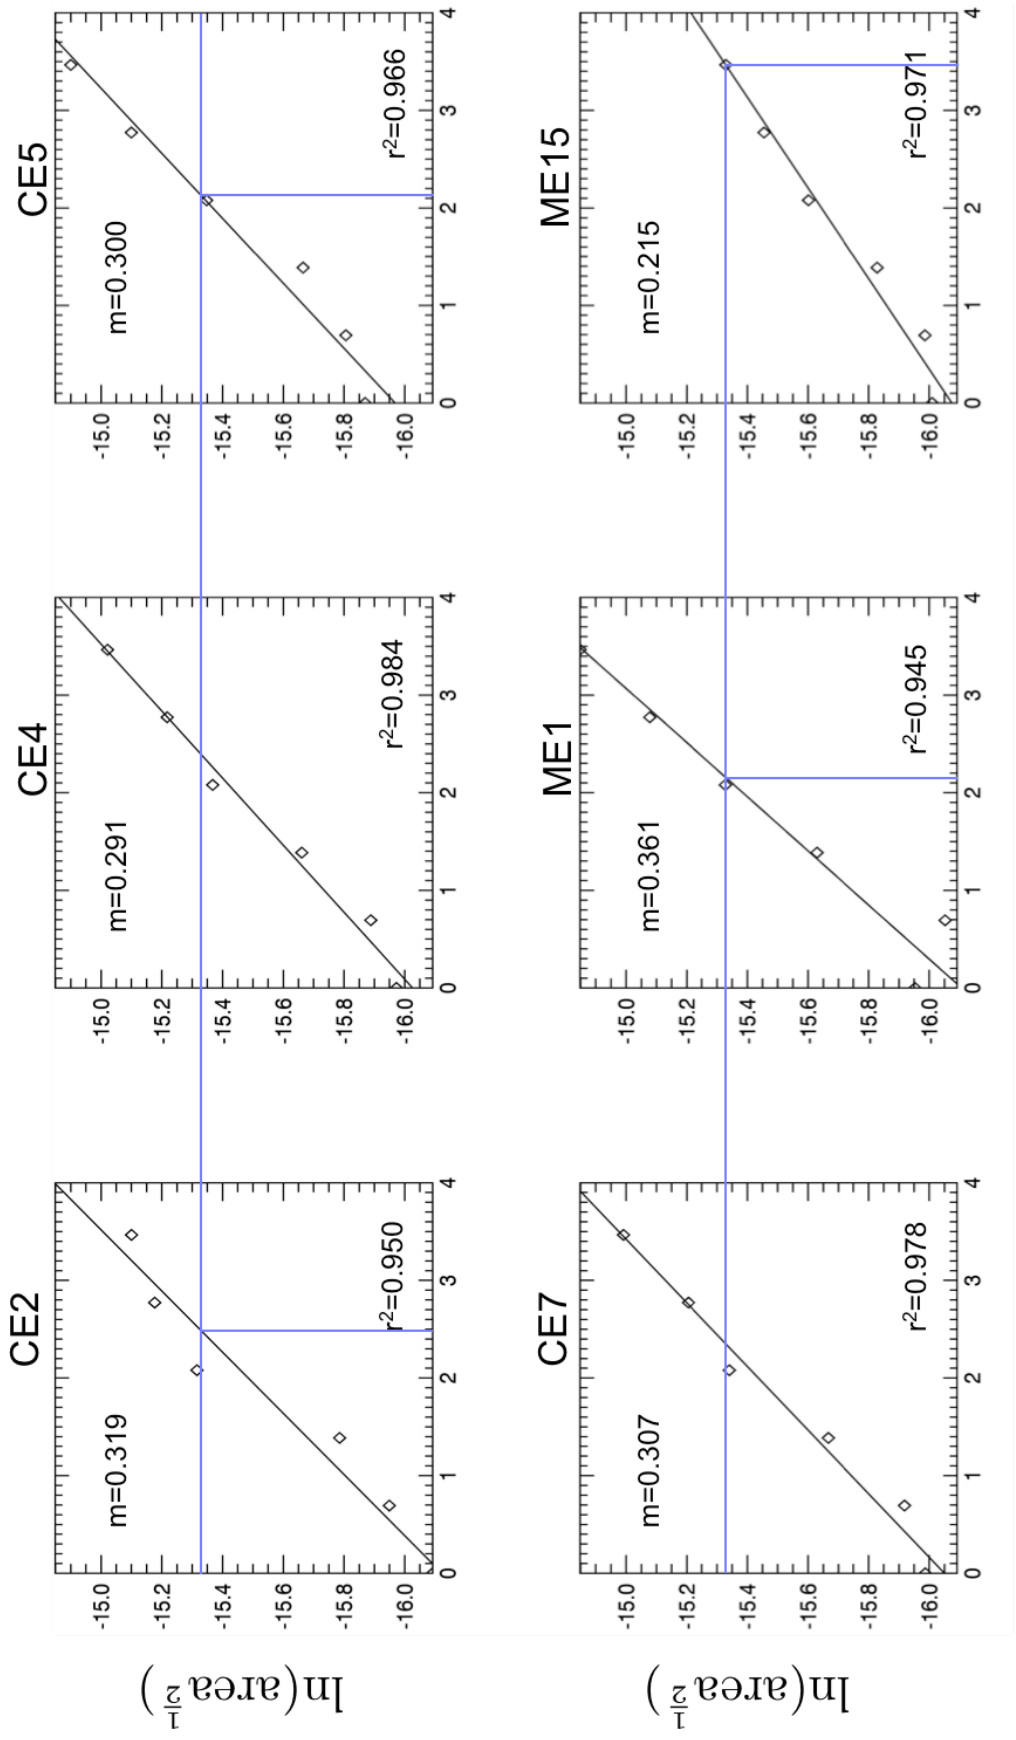

Figure 6.9: Trends in aging, with equivalent volume of precipitates plotted on the y-axis, and aging time in hours on the x. Strong linear relationships indicate adherence to LSW growth. Blue lines indicate selection of heat treatment times for creep testing.

Figure 6.9 shows the equivalent volume of precipitates, plotted against aging time. Lifshitz, Slyozov and Wagner (LSW) theory predicts a linear relationship between these values [217]. A best fit line is plotted for each alloy accompanied by the correlation coefficient. It is clear that these best fit lines closely fit the data. In addition, blue horizontal and vertical lines indicate how heat treatments were chosen for the alloys subjected to creep tests. Where these lines intersect the x-axis indicates the time to heat treat the alloy, with the goal of precipitates of the same size, matching where the blue lines intersect the y-axis.

6.3.2 Creep results

Micrographs for the four selected alloys are shown in Figures 6.10. The width of all of these images is $15\mu m$. The most obvious difference between these alloys is that alloy ME1 appears to have precipitates which have cross sections which are much closer to square than the other alloys. In addition, the precipitates of ME15 are considerably larger than those of the other alloys. This difference is likely explained by the air-cooling of the samples, which allowed for precipitate growth in these alloys that differed from that of the quenched samples.

SOMIMs and PMIMs for the four selected and heat treated alloys are shown in Figure 6.11 and Figure 6.12. These make the shape distributions more clear. Both the SOMIM and the PMIM of alloy ME1 have peaks that are close to the value of the square. The remaining alloys have peaks that are much closer to the value of the circle.

The kernel density estimates (KDEs) of Figure 6.13, calculated with an Epanechnikov kernel [218], display an estimated PDF of the data, essentially a smoothed histogram. The area under each sums to unity. Above each set of KDEs are four bars, centered on a diamond symbol. The diamond is the location of the overall mean for each alloy for the given shape feature. The bars extend to one standard deviation, positive and negative. The square symbol indicates the mean value for cuboidal shapes for the material. Vertical dashed lines indicate the median of the data.

Figure 6.14 shows the strain versus time curves for the creep tests conducted on the four alloys. A summary of the results of this test and the shapes of the γ' precipitates is given in Table 6.3. This includes the location of the overall mean shape values and the mean shape

Alloy	ω_1	ω_2	τ_1	τ_2	τ_3	γ'/γ	$\dot{\epsilon}_{\min}$ (s^{-1})	TTF (hrs)
CE-2 Mean	0.879	0.903	0.743	0.782	0.745	0.605	8.58e-08	157.9
Mode	0.966	0.974	0.919	0.921	0.914			
CE-5 Mean	0.898	0.925	0.775	0.819	0.793	0.611	2.02e-06	14.9
Mode	0.972	0.977	0.923	0.923	0.924			
ME-1 Mean	0.841	0.870	0.672	0.687	0.702	0.714	3.58e-11	414.4
Mode	0.943	0.931	0.859	0.769	0.834			
ME-15 Mean	0.844	0.862	0.684	0.707	0.673	0.746	1.59e-07	83.1
Mode	0.949	0.958	0.879	0.876	0.887			

Table 6.3: Summary of mean shape values and creep results for the examined alloys.

values for those precipitates with morphologies between those of the square and the circle. Area fraction is listed, as well as minimum creep rate, and time to failure (TTF) for each alloy in hours.

Figures 6.15, 6.16, and 6.17 show the results of linear regression for the natural log of the minimum creep rate of each alloy, and the mean, median, and mode of indicated shape features respectively. The mean values are accompanied by a set of bars indicating $\pm$ the standard error of the mean. The correlation coefficient is also shown. Between the figures, it is clear that the size of the particles has the lowest correlation to creep performance. The remaining features all show higher levels of correlation, suggesting that there may be a relationship. Overall, the median and mode show increasing levels of correlation. This may be explained by the increasing robustness of these values to outliers [219].

Figure 6.18 shows post-creep images of the four samples. The stress axis is normal to the image. The two alloys containing carbon, CE5 and ME15, are particularly difficult to image using BSEs. The contrast does not allow for good imaging of the precipitates. However, it is clear from these images that the two alloys which performed best in creep, CE2 and ME1 both exhibited dramatic shape evolution during creep, but in two divergent ways. CE2 shows a rafted structure while ME1 has precipitates that have coalesced into large irregular arrays.

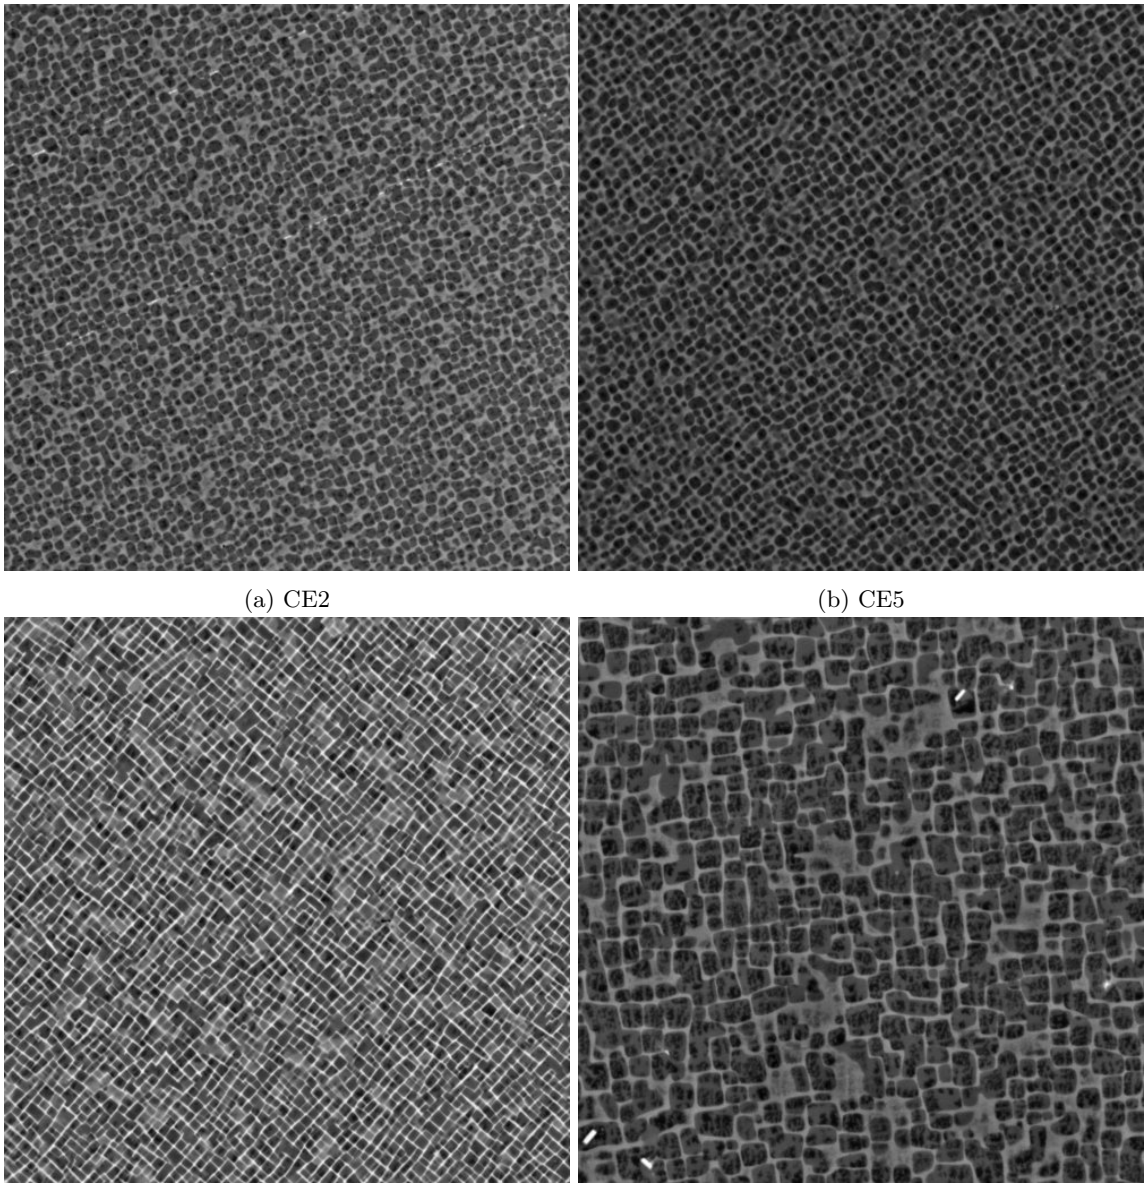

(a) CE2 (b) CE5

(c) ME1 (d) ME15

Figure 6.10: Micrographs of superalloy microstructures prior to creep testing. All images have a width of $15\mu m$.

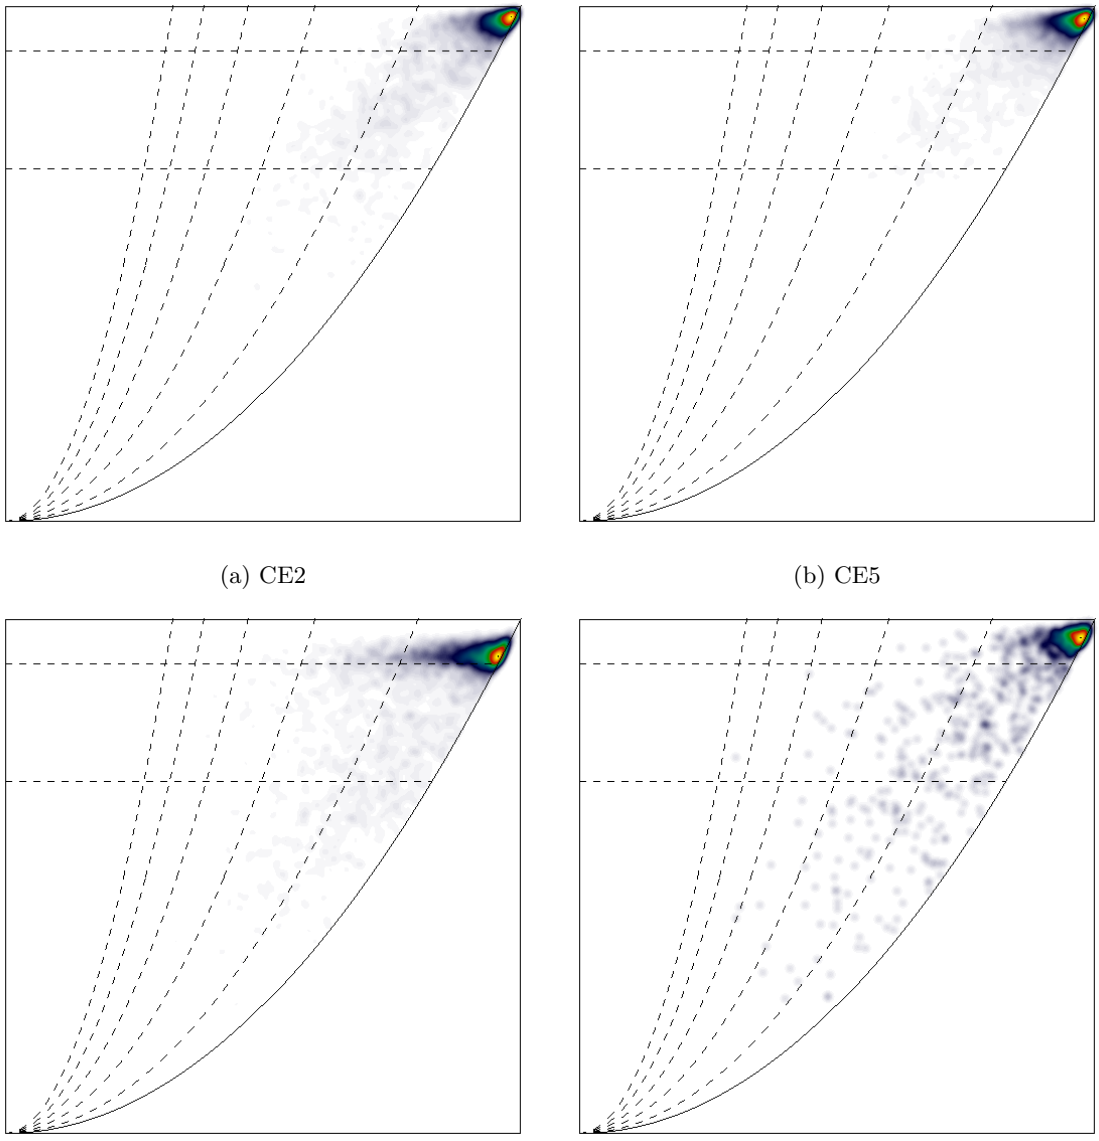

(a) CE2

(b) CE5

(c) ME1

(d) ME15

Figure 6.11: SOMIMs for creep-tested alloys.

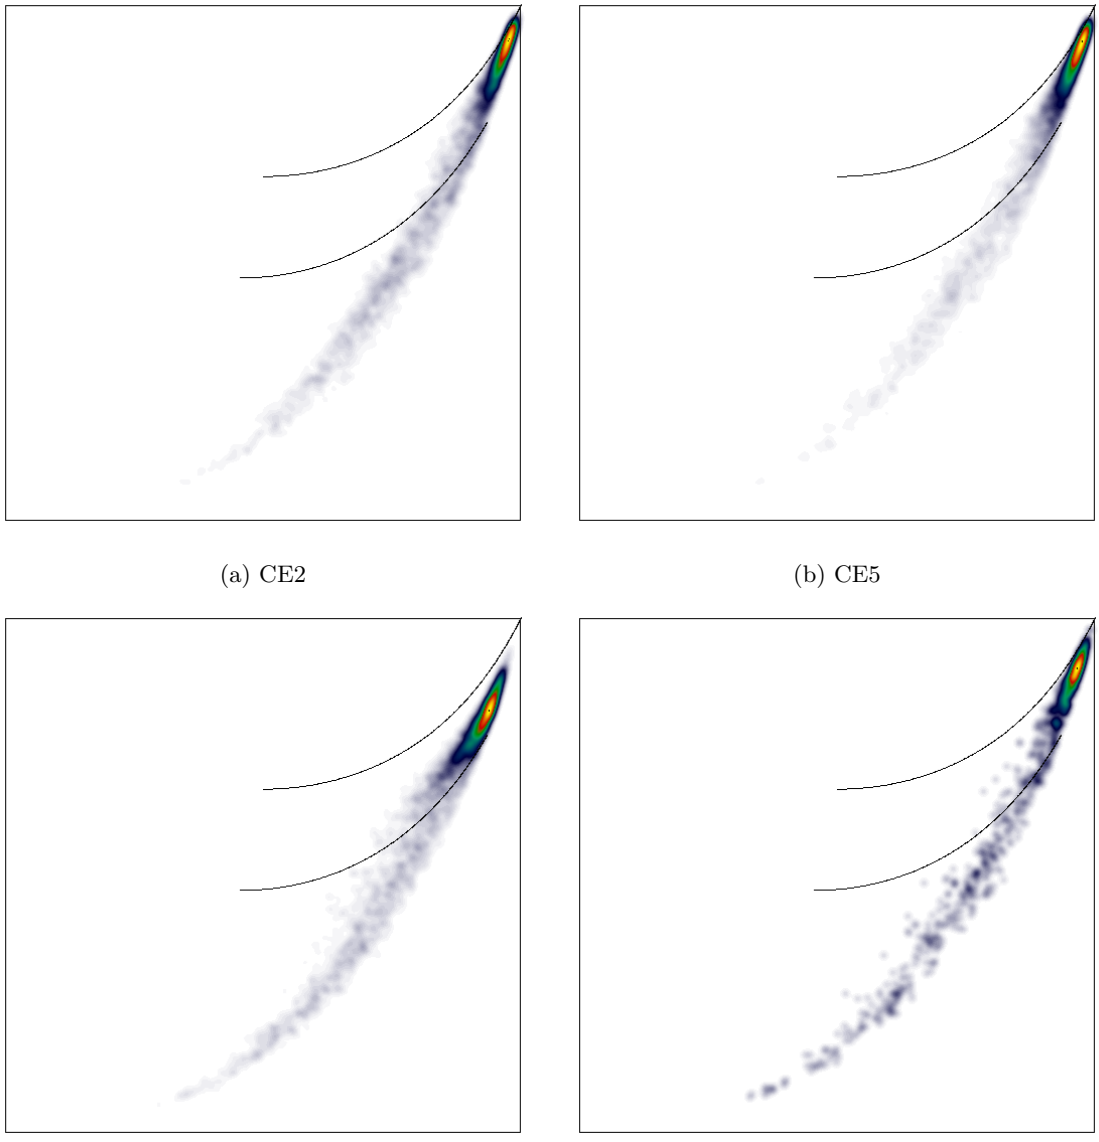

(a) CE2

(b) CE5

(c) ME1

(d) ME15

Figure 6.12: PMIMs for creep-tested alloys.

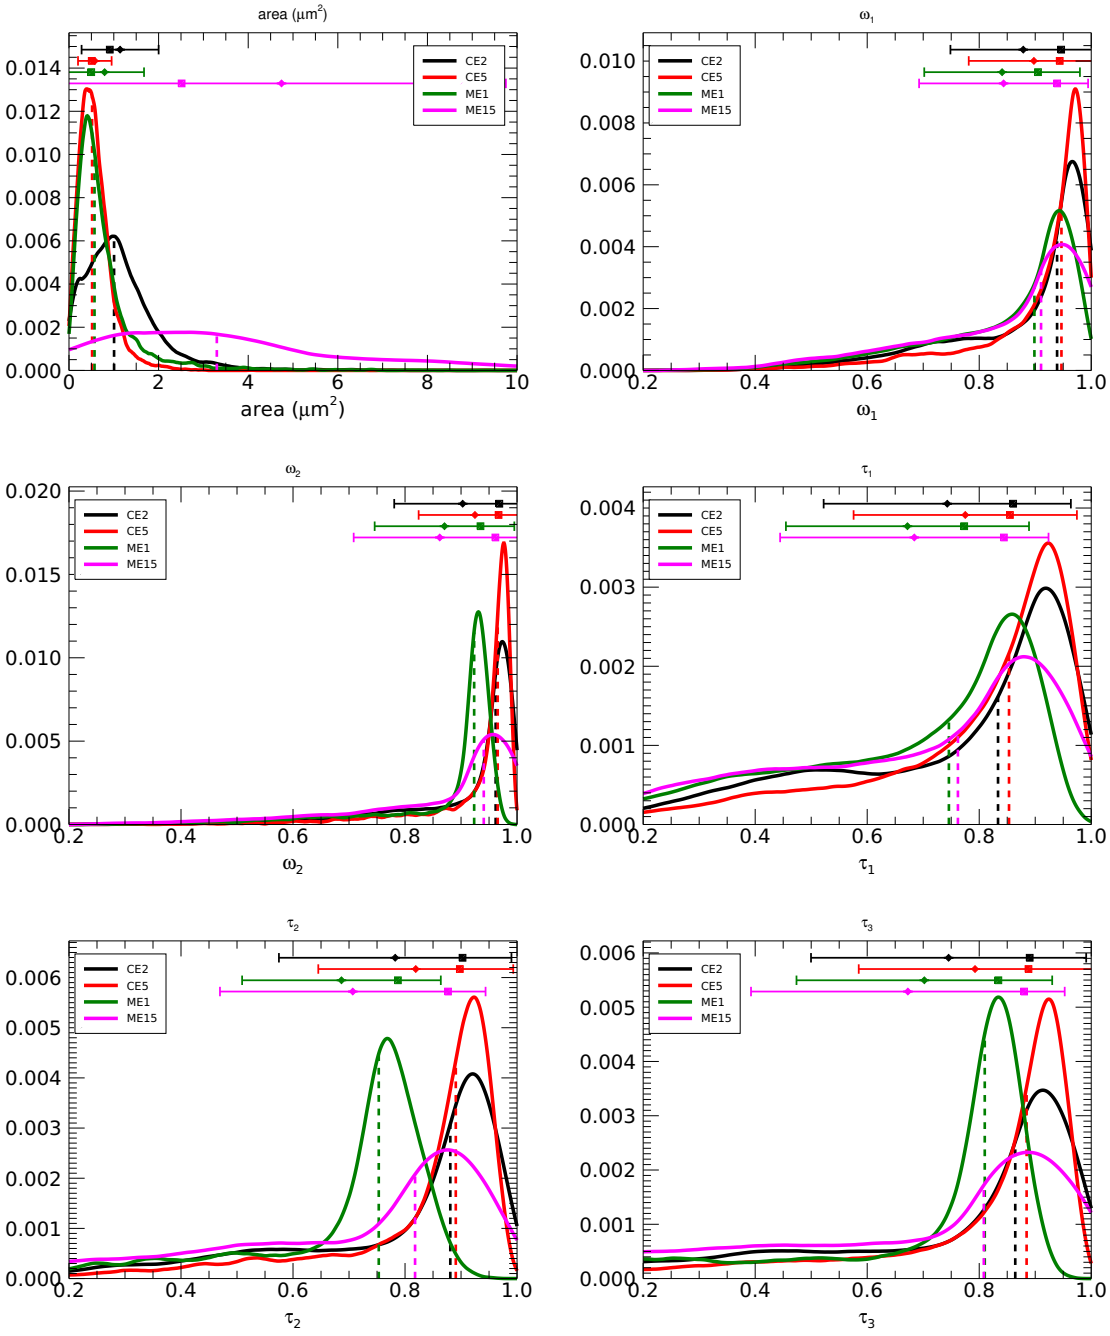

Figure 6.13: Kernel density estimates for the four examined materials and the indicated properties. Bars above the plot indicate overall mean as a diamond, plus or minus the standard deviation of the data. The square symbols indicate the position of the mean of shapes which have a cuboidal morphology. Vertical dashed lines indicate the median of the data.

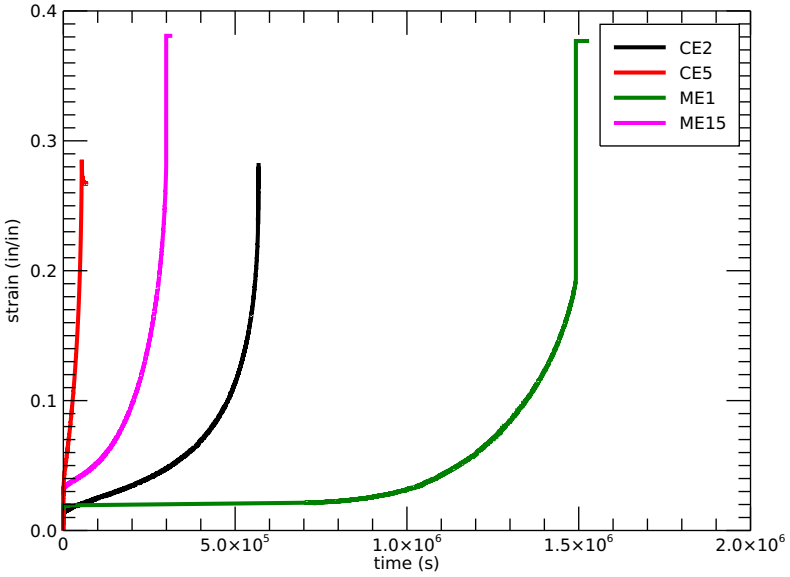

Figure 6.14: Stress versus time in seconds for the alloys examined here. ME1 showed the greatest time to failure and overall creep resistance while CE5 showed the least.

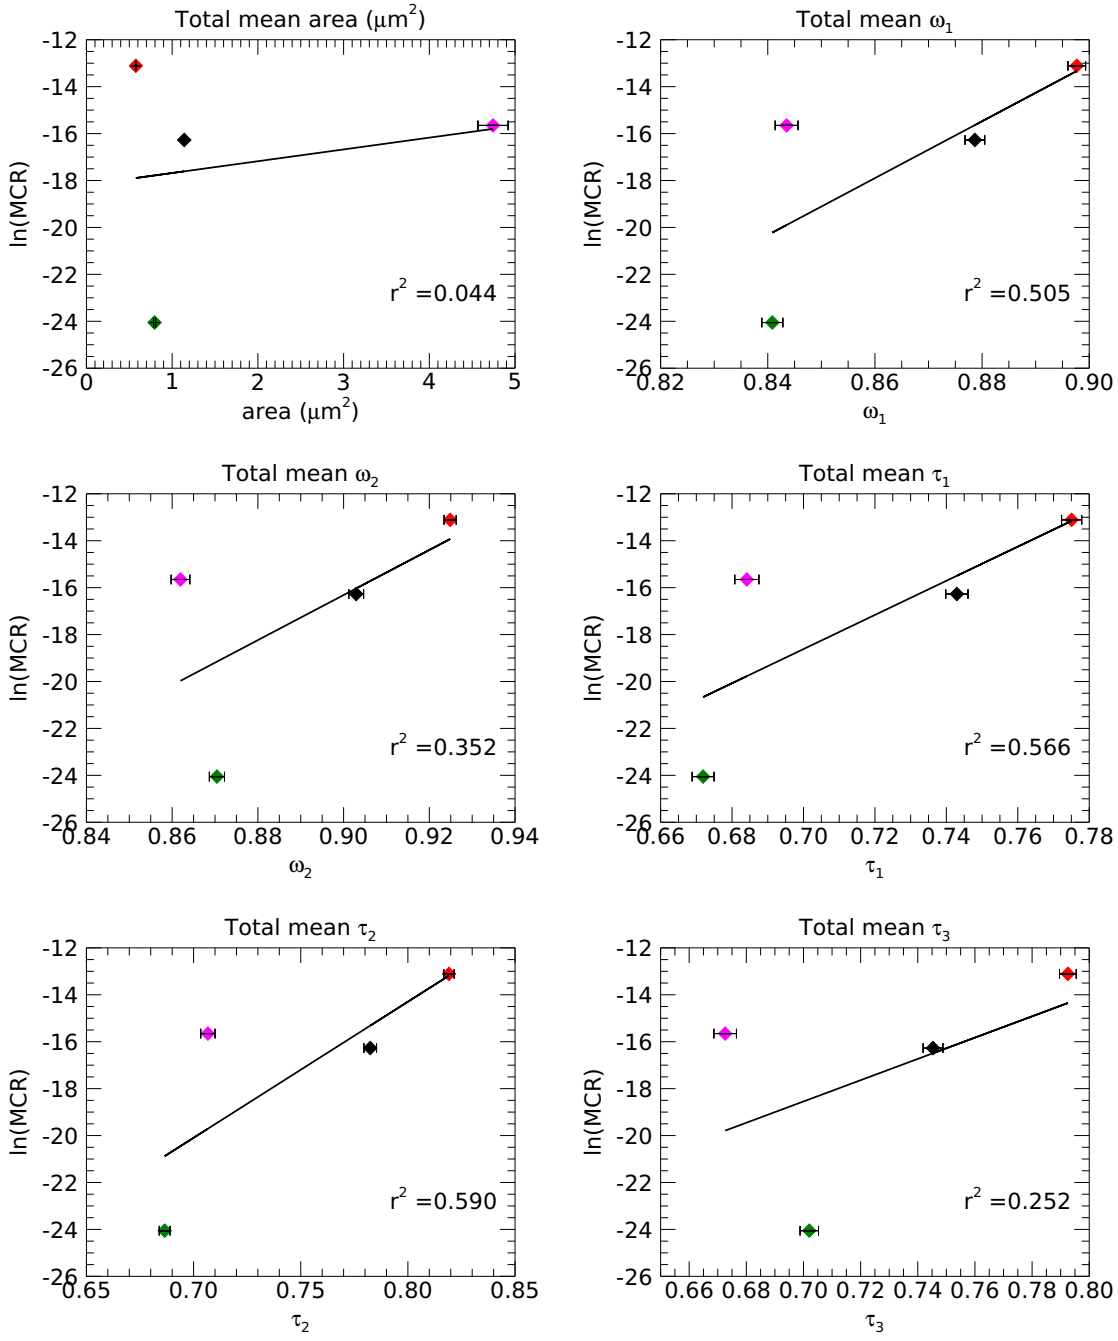

Figure 6.15: Best fit line for overall mean of the features plotted against the natural logarithm of the minimum creep rate for each alloy. Colors of symbols indicate the alloy, matching those of Figure 6.14.

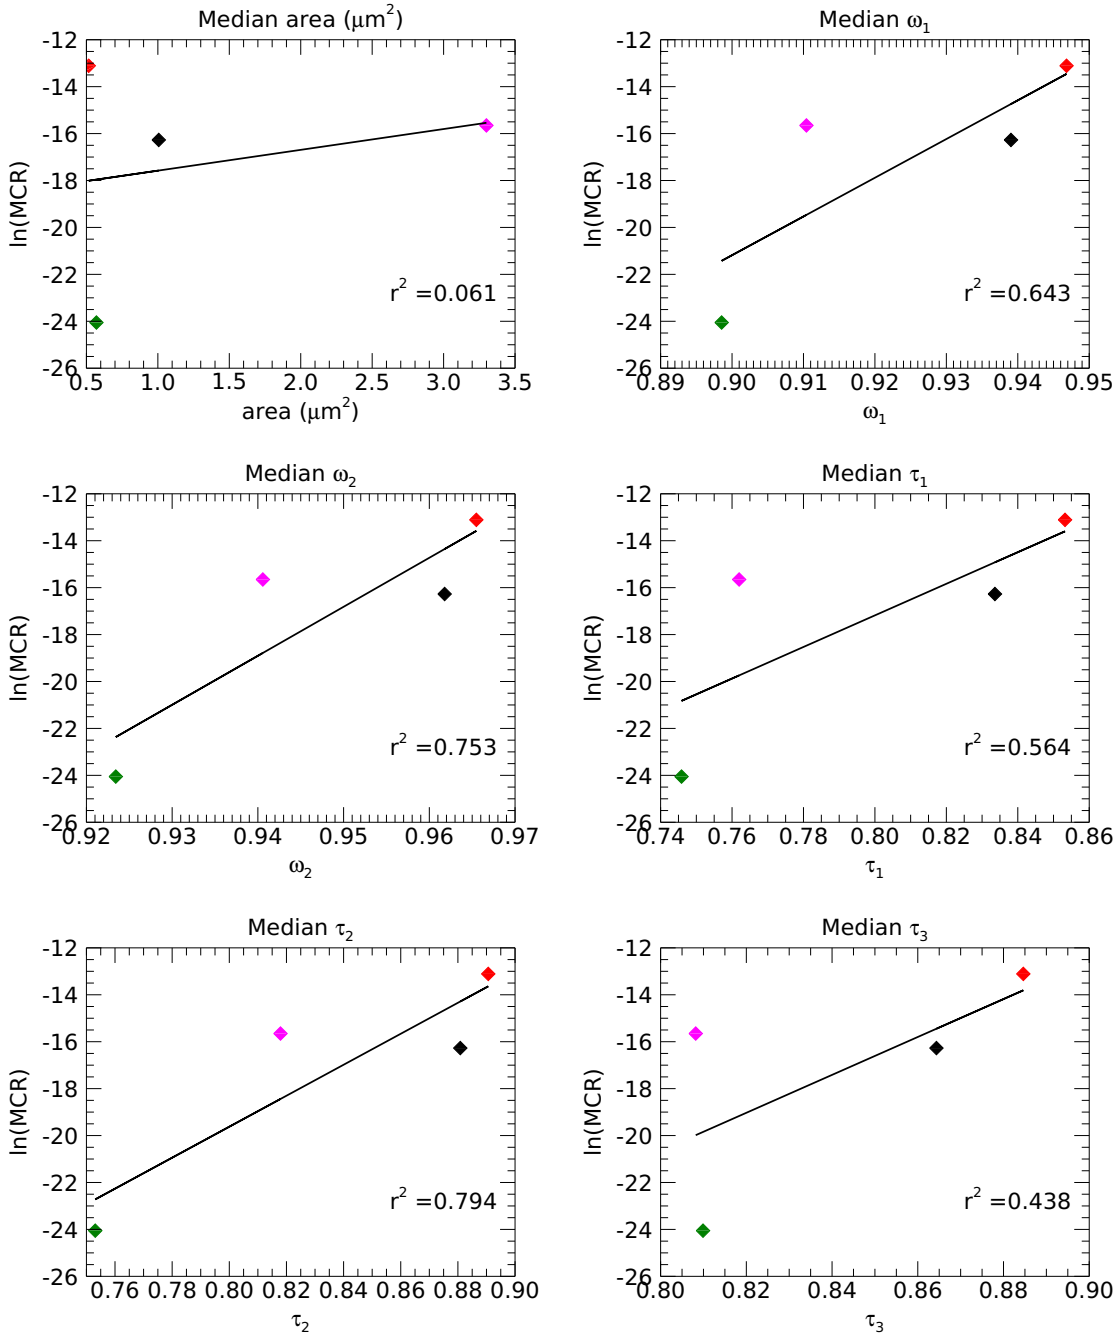

Figure 6.16: Best fit line for the median values of the features plotted against the natural logarithm of the minimum creep rate for each alloy. Colors of symbols indicate the alloy, matching those of Figure 6.14.

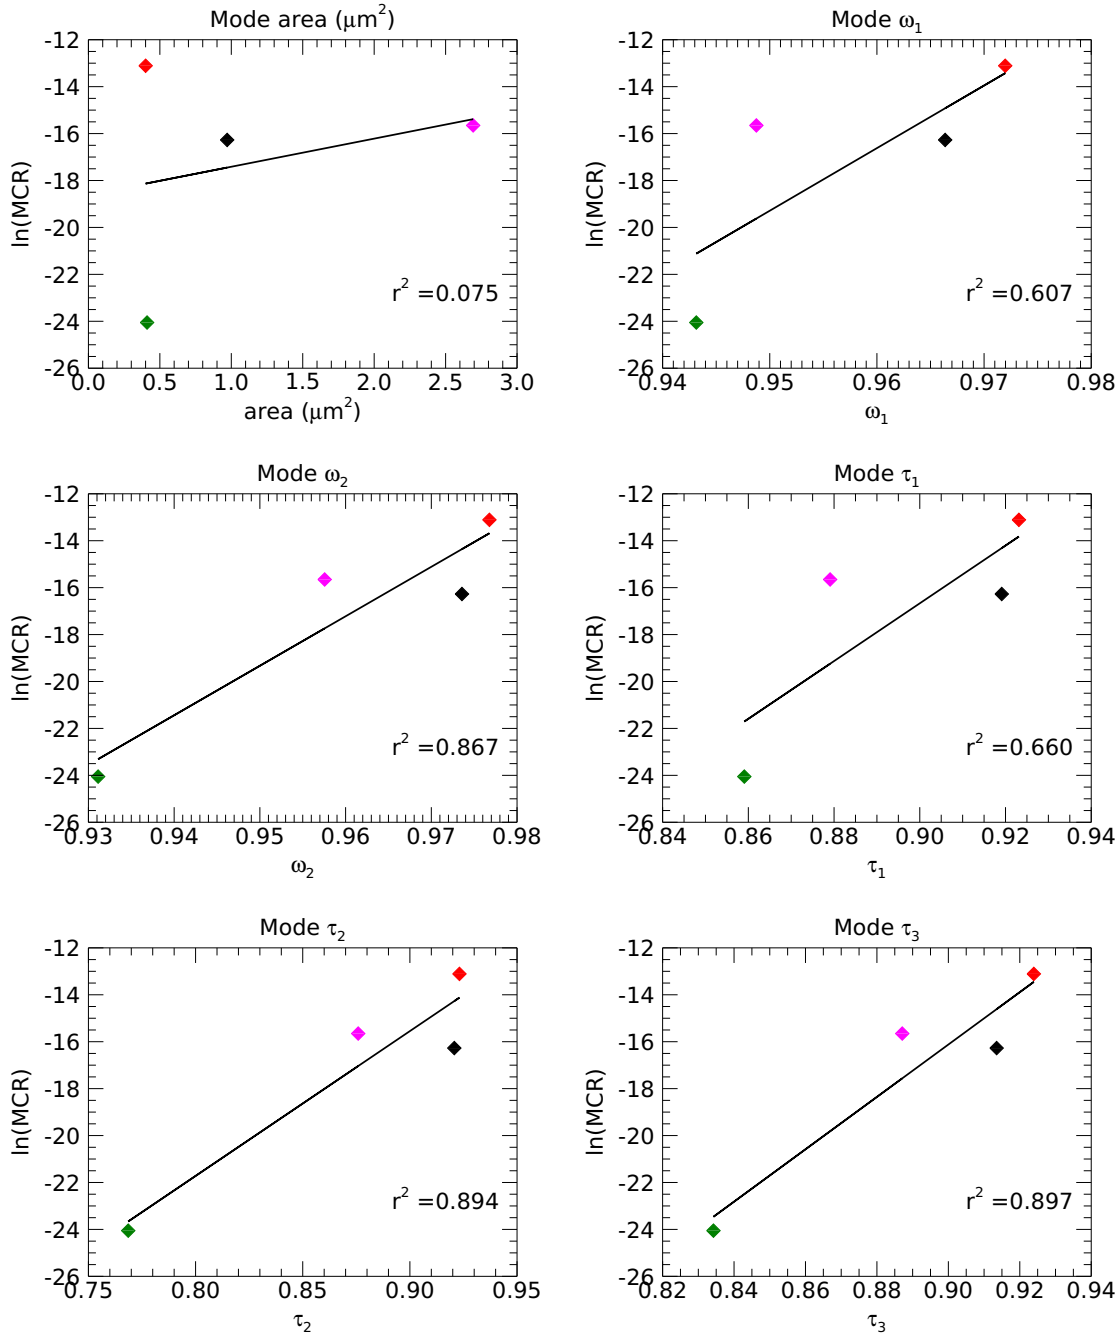

Figure 6.17: Best fit line for the mode values, taken from KDEs of Figure 6.13 of the features plotted against the natural logarithm of the minimum creep rate for each alloy. Colors of symbols indicate the alloy, matching those of Figure 6.14.

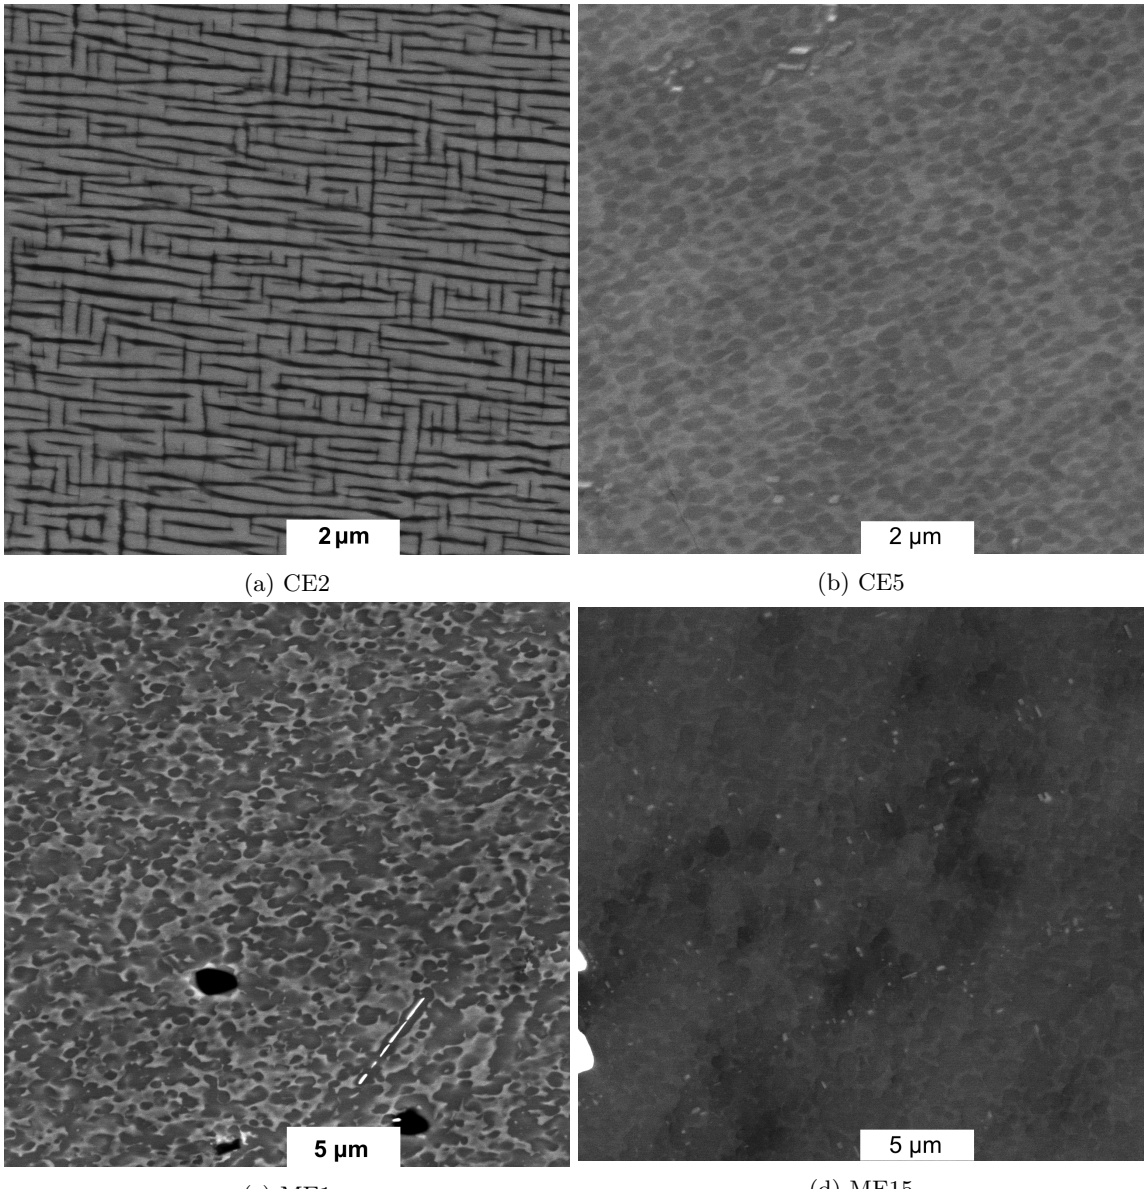

(a) CE2

(b) CE5

(c) ME1

(d) ME15

Figure 6.18: Post-creep micrographs. Stress axis is normal to the images. Lighter phase is γ'.

6.4 Discussion

6.4.1 Aging

Analysis of the combination of SOMIM maps and individual ω plots reveals the nature of the shape evolution with aging. The SOMIMs of Figure 6.6 and the shape descriptor plots of Figures 6.7 and 6.8 support the idea that at low aging times the precipitates tend to have more irregular shapes and larger aspect ratios. This is indicated by a larger spread of the peaks in the maps, and the relatively lower values of both ω_1 and ω_2 in the ω/time plots.

This kind of behavior is not expected. As discussed in Section 6.1.1, it is generally understood that γ' precipitates in these materials begin as spherical particles with low volumes, and only after some growth begin to display the more cuboidal shapes which minimize elastic energy. Prikhadko and Ardell observed similar behavior when tracking their shape descriptor, Σ, against precipitate size [209]. It was noted that as particles grow, they become more circular, followed by a slow return to a less circular shape parameter. This period of increasing regularity and decreasing aspect ratio is observed in Figures 6.7 and 6.8 for most alloys for treatment times < 10 hours.

It is possible that these effects are due to difficulties in the segmentation of SEM images at extremely high magnification. In the case of these materials, at one hour of aging, the distance between precipitates can be as low as 20nm, which borders on the resolution limit for SEMs. This could lead to incorrectly merged precipitates during segmentation, which would increase the aspect ratio and the irregularity of the measured shapes and lower the average values of both ω descriptors. At higher aging times and precipitate sizes, this effect disappears and the precipitates assume the more expected behavior.

Van Sluytman and Pollock [211] and Prikhodko and Ardell [209] both observed a slowing in the change of shape class (square to circle) with aging time and the approach of precipitates to some final shape. This same effect is visible in the alloys examined here. Once through the initial period of rapid precipitate shape change, the alloys here appear to hit a region of constant precipitate morphology. This is most obvious from the plots of Figure 6.8. Every alloy, with the exception of perhaps ME1, appears to have a very slow moving or perhaps unchanging ω_2 value after approximately 16 hours of aging. This is despite the

fact that the precipitates are continuing to grow in size as seen in Figure 6.9.

All of this indicates that, in the absence of applied stress and at the aging times examined here, the precipitates quickly assume a stable shape and aspect ratio and maintain it for extended lengths of aging, despite continued precipitate growth. It would require continued heat treatments in order to see if the precipitates in these alloys will continue this trend or at some point enter another region of morphological change.

Figure 6.7, showing trends in ω_1 and therefore aspect ratio, shows almost no discernible relationship between its shape parameter, ω_1 and aging time. This agrees with the work of Prikhadko and Ardell who found no relationship in unstressed alloys, but is in contrast to that of Fährmann *et al.*, who found increasing aspect ratios in higher misfit alloys with aging [208]. The reasons for this behavior are not clear. In unstressed alloys, precipitate coalescence, which is expected with continued precipitate growth, cannot be predicted to favor any one direction. However, because ω_1 is rotationally invariant, *any* increase in aspect ratio, regardless of direction, would contribute to a decrease in this value. It is possible that, because no direction is favored, a precipitates will be likely to merge in two orthogonal directions at once, which maintains aspect ratio but has the potential to lower ω_2.

Figure 6.9 was created primarily as a check of the segmentation quality. LSW theory predicts a linear relationship, which is confirmed by the quality of the match of the best fit lines. The major differences between the alloys is the slope of this line, which is typically described in the literature as a function of volume fraction. The average volume fraction for each alloy, calculated from the segmented images, is as follows: CE2: 57.0%, CE4: 53.5%, CE5: 51.1%, CE7: 53.3%, ME1: 68.6%, ME15: 69.10%. There is no obvious relationship here, as the alloy with the largest slope: ME1, and the alloy with the smallest, ME15, have very similar volume fractions. In addition, at low aging times, one may observe that there are greater deviations from the best fit lines, supporting the idea that there are segmentation problems at lower aging times.

6.4.2 Creep and precipitate morphology

Size differences

It is clear that the erroneous introduction of air-cooling over water quenching affected the samples to some degree. Example micrographs in Figure 6.10 were all taken at the same magnification, with image widths of $15\mu m$. Alloy ME15 has precipitates that are considerably larger than the others. Kernel density estimations (KDEs) of the size of precipitates, shown in Figure 6.13, clearly show that the distribution of the size of γ' particles in ME15 does not only have a significantly higher mean value, but is also dramatically wider as well. The width of the distribution suggests that some of the precipitates of the ME15 alloy have either begun to coalesce, or to develop the more irregular particle morphologies associated with a loss of coherence [12].

In contrast to ME15, the remaining alloys had much more similar precipitates in terms of cross-sectional area. Again, from Figure 6.13, both CE5 and ME1 show very similar, almost identical, size distributions, with very similar mean and median values. CE5 has a larger spread than these alloys and a larger overall particle size, with a larger mean and median size as well.

ME15 shows such a dramatically different range of precipitate sizes, most likely due to the air-cooling of the samples, rather than quenching, as should have been done. It is interesting to note that despite the change in heat treatment cooling rates, the other three alloys display similar sizes but ME15 does not. This is probably caused by compositional differences, and the impact of the longer aging time for this alloy combined with the change in heat treatment. Regardless, as will be discussed below, it does not appear that precipitate size played a large role in influencing the creep resistances of these materials, at least in this study.

Precipitate shapes

The SOMIMs and PMIMs of Figure 6.11 and 6.12 for the creep-tested alloys highlight what is suggested by the micrographs of Figure 6.10. The precipitates of ME1 display a morphology is that is significantly different than the other alloys. The precipitates appear much more square, and this is supported by the SOMIM, which has a peak right above the

location of the square. Both CE2 and CE5 have precipitates that are closer to circular. The precipitates of ME15 appear to be somewhere between those of the CE alloys and ME1, however, the SOMIMs and PMIMs make it clear that the morphology of these particles is closer to the CE alloys. The tail of the distribution of ME15 also has a more spotted appearance. This is a result of the precipitate sizes, which reduced the number of particles that were measured from roughly 5000 for the other alloys to 1000.

The kernel density estimates (KDEs) of Figure 6.13 show essentially smoothed histograms of the data, normalized so that the area under the curve is equal to unity. The bars along the top indicate ± one standard deviation from the mean, indicated by a diamond shape. The square shape indicates the mean of the indicated shape feature in each alloy for particles that have ω_2 values greater than or equal to 0.91. This allows for the quantification of only the regular shapes which have morphologies between the circle and the square. Dotted vertical lines indicate the median of the data, and the peak of each KDE is the mode.

It is clear from the shape of the KDEs that ME15 has the largest spread for each shape feature. This is suggested by the fact that ME15 has the lowest peak in each KDE, indicating more particles are farther away from the mode than in any other alloy. This is particularly true looking at ω_2, τ_2 and τ_3. This suggests that ME15 has many more irregular shapes and is therefore further along in particle growth than the other alloys; its precipitates either beginning to coalesce or lose coherency, which occurred prior to creep testing, reducing its creep properties.

Another notable feature of these plots is the difference between the means, medians and modes of these data sets. The long-tailed skewed distributions that are displayed indicate the presence of many outliers. Out of the three measurements, the mean is most sensitive to these, followed by median, and then mode [220]. When looking at mean values, the alloys tend to be grouped by their type: ME alloys nearest to each other and CE alloys together. This is true in the case of all the shape descriptors, with the exception of τ_3, which has a more evenly distributed set of means.

Looking at medians, this is only the case with ω_1, τ_1, and τ_3. The medians of the ω_2, and τ_2 values are more evenly spread. The modes of these values, indicated by the peaks

of the KDEs are either evenly spaced over a small range, as in the case of ω_1 or τ_1, or distributed in such a way that the mode of ME1 is clearly separate, as in the case of τ_2 and τ_3.

The implications of parameter choice are important and should reflect the expected behavior of the properties being examined. If the relevant property, in this case creep resistance, is expected to be influenced only by the most common particles, then mode should be used. Mean may be a good choice if outliers are expected to have an effect, but it may be that the mean value itself is influenced by the outliers to a larger degree than the affected property. It is unclear what is the case with these particular alloys. Regression analyses are therefore presented using all three parameters.

Creep resistance

The strain-time curves presented in Figure 6.14 show the results of the creep tests on these alloys. The best performing alloy was ME1, followed by CE2, ME15, and CE5. The quantification of creep resistance was done by calculating the minimum creep rate, the results of which are shown in Table 6.3. Minimum creep rate was calculated by fitting a series of functions of the form $y = ax^b + c$ to each line. The minimum creep rate was taken as the minimum slope of these fitted curves. Table 6.3 also presents the time to failure in hours as an alternative measure, but these values were not used here. Best fit lines for the shape features and natural log of the creep resistances of each alloy are presented in Figures 6.15, 6.16, and 6.17 for the mean, median and modes of the data, respectively.

Looking at the mean plots of Figure 6.15, there appears to be some relation between some shape features and the creep resistance. The relationship is strongest in the cases of ω_1, τ_1, and τ_2. These all display correlation coefficients greater than 0.5, indicating that over half of the variance in creep resistance can be attributed to variance in these shape values. However, the correlation coefficient of ω_2 is low at 0.35. Previously mentioned work by Van Sluytman *et al.* [211] utilized values directly related to ω_2 and found indications of correlation.

With the use of median values, the strength of *all* the correlations increases. The r^2 values of ω_2 increases to 0.75, and the same value for τ_2 increases to nearly 0.8. These

values indicate a strong relationship, and suggest that the presence of outliers is masking the overall effect of shape upon creep resistance. Finally, using the mode of the shape features increases the correlations once again to 0.87 in the case of ω_2 and nearly 0.9 in the case of τ_1 and τ_2.

The fact that, even when changing the quantifier of the distribution from mean to median to mode, the correlation between creep resistance and size did not grow significantly suggests that the increased correlations are not the result of simple data manipulation, but instead are revealing relationships between these properties which were disguised by the presence of outliers.

The improved correlation with increasingly more robust distribution measures suggests that the outlier particles, whose influence is being reduced, have less of an effect on the overall creep resistance than the shape of the majority of the particles represented by the mode. This is compatible with what is known regarding high temperature creep deformation. Dislocations travel through the matrix and form networks which are restricted to the surface of the precipitates. Rowland noted that spherical precipitates do not form interfacial dislocation networks [221]. In addition, less cubical particles with lower lattice misfits do not create rafts which inhibit dislocation motion. The question may be analogous to percolation. If the majority of the particles are of the ideal morphology then they will inhibit dislocation motion and improve creep resistance. If this is the case, then the mode of the data is the proper quantifier. However, to fully explore this, much more analysis is necessary, including imaging and testing of more samples with a variety of shapes and shape distributions.

Some care should be taken in trying to draw too many conclusions from this data. For example, lower ω_1 values are correlated, up to an r^2 value of 0.6, with improved creep resistance. This would indicate that increased γ' aspect ratios improve the performance of the material. However the trend is primarily drawn between ME1 and the CE alloys. While it is true that ME1 has particles with larger aspect ratios, the alloy also has precipitates of a much more rectangular nature which is a completely separate quality that is measured by ω_2. This correlation between ω values in the samples examined here may mask underlying effects, or suggest effects that are not present. There are not enough data points to make

distinctions between these effects as of yet, but with additional tests of more alloys, more subtle details can be uncovered.

Additionally, there is likely a large degree of other correlations between these shape descriptors. It is possible that for the classes of shapes that are present in these alloys, many of the τ values are so strongly correlated with the ω values that measuring the shapes to the increased degree of detail provided by the fourth order τ values is unnecessary. However, at this point in time, without looking at more samples, it is difficult to know for sure. Regardless, as previously described in Section 6.1.2, interactions between particle morphology and creep resistance has been observed in the past, so these observations are consistent with previously noted qualitative relationships.

The correlations that are described here are all based upon an assumption of a linear relationship between the shape features and the creep resistances of these alloys. While there are strong correlations present, there is no reason why the relationships suggested by this data should necessarily be linear. A larger number of alloys with a wider variety of precipitate shapes would allow for a more detailed and complete analysis.

Looking at the micrographs of Figure 6.18, it is clear that the best performing alloys, ME1 and CE2, have undergone a significant degree of precipitate evolution, while the other two alloys have not. This may indicate, as discussed in Section 6.1.1, that both ME1 and CE2 have a higher degree of misfit at the creep temperature than the other alloys. This increased misfit would drive rafting and contribute to the creep resistances of the materials.

The variations in particle morphology between ME1 and CE2 may be due to a difference in the sign of the lattice misfit. The stress axis is perpendicular to the images of Figure 6.18. CE2 has therefore formed rafts parallel to this axis, indicating positive misfit. It is difficult to tell for ME1, but the extended conglomerations of γ' strongly indicate rafts which have formed close to the plane of the image, similar to microstructures shown elsewhere [205]. This indicates perpendicular rafting, suggesting negative misfit. These rafts are thought to inhibit dislocation motion and contribute to strengthening of the alloy in creep. This is consistent with the superior creep resistance of this alloy.

From all of the plots, it is clear that the ME15 alloy is the most significant outlier and therefore has a lower creep resistance than would be suggested by the shape features

present. The most obvious difference between ME15 and the other alloys is the large size of the precipitates. Creep resistance *is* affected by the size of precipitates. Nie *et al.* used finite element analysis to predict such an effect in similar superalloy materials and then observed corroborating evidence from experiments [222]. However, Figures 6.15, 6.16, and 6.17 all show very little correlation, suggesting that the influence of size between these alloys is small. Statistically speaking, by using four alloys, three of which with similar precipitate volumes, the effect could be masked. Further examination of a wider array of alloys with an increased variation in precipitate sizes would allow for more subtle effects to be noted.

6.5 Conclusions and future work

The work described in this chapter provides a stepping stone toward fully understanding the influence of precipitate shape on the creep resistances of single-crystal nickel-base superalloys. While relationships had been noted in the base, they were primarily based upon qualitative observations of shape. This work utilizes quantitative methods to describe the shape of precipitates to a high degree of detail. In addition, examines multiple aspects of γ' precipitate shape separately.

Prior to this work, it was unclear if automated methods would provide the quality of segmentation and analysis necessary to allow for an accurate measurement of precipitate morphology. The above results demonstrate the effectiveness of the current approach in measuring superalloy precipitates sizes and shapes utilizing an automated method relying on Cartesian moment invariants as the shape descriptors. Size measurements closely match the LSW model of growth. Shape measurements indicate no obvious effect of aging on aspect ratio, except for a mild effect on those alloys with precipitates of more square cross-section. Precipitate shapes tend to start more irregular or square, and move toward more regular, circular cross-sections up until about four to eight hours of aging. This transition does not match expected behavior but may be due to imaging difficulties at small sizes. After this region, the behavior does match what is expected.

While the distributions of shapes were strongly affected by outliers, this is mitigated to some extent through using more robust quantifiers of these shape distributions, such as the medians or modes of the data sets. Ultimately, strong correlations between ω_2, τ_2,

and τ_3 with the minimum creep rate were shown. However, the number of alloys examined do not allow for the full extent of the relationship to be made clear. Nevertheless, this work supports the hypothesis, and shows that moment invariants can be an effective tool in drawing quantitative relationships between microstructural shapes and material properties in single-crystal nickel-base superalloys.

6.5.1 Future work

Continued examination of these materials and the shapes of their precipitates should focus on extending the analysis to a greater number of alloys to improve the statistical support for the presence of a relationship, and to more fully explore the shape of this relationship as well.

CHAPTER VII

Summary and Conclusions

7.1 Summary

The goal of this work is to explore the relationship between microstructural shapes and the properties or function of the material. Using advanced segmentation methods, and high-dimensional visualization and clustering techniques, this analysis can be performed across a wide variety of materials and on large numbers of shapes automatically. This type of characterization is an under-explored aspect of materials science, and quantitative correlations between morphology and material performance serves to illustrate that showing that these types of links exist and are meaningful while simultaneously illustrating the effectiveness of the techniques described here.

The hypothesis of this work is as follows: *Shapes of microstructural subregions in embryonic* Arabidopsis Thaliana *leaves, additive manufacturing feedstock powders, and single-crystal nickel-base superalloys, as quantified by Cartesian and Zernike moment invariants, correlate to the overall properties or function of the material as a whole.* This hypothesis was explored in the following material systems under the given conditions:

1. The three-dimensional shapes of cells in *Arabidopsis Thaliana* cotyledons used to identify cell function.

2. The two-dimensional shapes of a subset of additive manufacturing powder feedstock used to identify the material and manufacturer of the feedstock.

3. The two-dimensional shapes of γ' precipitates in single crystal nickel-base superalloys correlated to the creep resistance of the alloy.

In the case of the *Arabidopsis Thaliana* cotyledon cells, this work has shown the efficacy of the segmentation, dimensionality reduction, and clustering methods of the unsupervised machine learning methods described previously, as well as the success of a trainable support vector machine classifier. The best success rate utilizing unsupervised methods has a mutual information value of 0.627, and based upon qualitative analyses of clusterings is it clear that 3D moment invariants quantify relevant shape information of these cells which correlate to cell role and function. The use of unsupervised methods show that cells of similar function group naturally in the shape data.

Using supervised methods, a maximum success rate of above 90% was achievable, but this was done utilizing features which include shape, volume, nearest neighbor count, and orientation. When restricted to shape alone, success rates of 60-70% were achievable. The success of supervised methods show that these features can be used to artificially separate the classes of cell.

Work with superalloy feedstock has shown that two-dimensional Cartesian moment invariants are capable of quantifying subtle differences between particles. Utilizing a bag of words approach to image classification allows for a success rate of 90% in the identification of 80 test images pulled from a collection of images of eight different material powders. Using five lower order moment invariants allows for nearly the same success rate of approximately 86%. Initially, work which focused on simple histogram matching proved ineffective, but the use of a bag of words model, which allowed for greater influence of irregular shapes which made up the minority of the particles, improved matching significantly. In addition, t-distributed stochastic neighbor embedding (t-SNE), provides a way to reduce the dimensionality of the data while still retaining a large degree of structure. This algorithm, coupled with a density based clustering method, allows for the clustering of circular and near circular particles by subtle morphological features which also improves clustering.

In the case of the superalloys, the work here has shown that there is correlation between several shape quantifiers and the creep resistances of four experimental superalloys as measured by the minimum creep rate. It appears that with regard to all five shape parameters examined (ω_1, ω_2, τ_1, τ_2, τ_3), reduced values, indicating values closer to square, correlate with improved performance in high temperature creep. This matches qualitative

observations, but there are a number of caveats. The first is that with only four examined alloys, it is difficult to fully examine the full nature of the relationship. Correlations shown in this work are linear but there may be a more complex interaction between shape values and creep resistance. Also, with only four examined alloys, correlations between shape values may be masking underlying effects. Finally, without post-creep imaging, it is difficult to determine the true nature of dislocation-precipitate interactions during testing.

7.2 Conclusions

It is clear that in the three systems examined here that the shape of microstructural subregions either reflects material function and properties, or influences them directly. Classifying shape is therefore a powerful tool when it comes to understanding the relationships between material structure and material function. Moment invariants provide a way to quantify shape that is complete, able to be tuned to an arbitrary level of detail, and capable of quantifying information that is relevant to the materials scientist.

The three material systems examined here are not the only ones in which shape is an important factor. Extending the methods described here to other systems only gets easier with time. Methods for segmentation are becoming more and more effective, and improvements to machine learning algorithms allow more interesting and subtle relationships between features to be examined. Hopefully, shape will become an increasingly investigated property of material microstructures in the future.

In closing, this quote from Dr. Cyril Stanley Smith is particularly applicable and sums up the motivation for four years of work spent "examining the background":

> "As science matures, attention is focused more on forces, and shape tends to be taken for granted... Nevertheless, it is desirable from time to time to examine the background and to attempt to deduce relationships of greater complexity and utility than those that are immediately obvious" [112].

BIBLIOGRAPHY

[1] Nesma T Aboulkhair, Nicola M Everitt, Ian Ashcroft, and Chris Tuck. Reducing porosity in alsi10mg parts processed by selective laser melting. *Additive Manufacturing*, 1:77–86, 2014.

[2] Herbert Freeman. On the encoding of arbitrary geometric configurations. *IRE Transactions on Electronic Computers*, (2):260–268, 1961.

[3] William I Grosky and Rajiv Mehrotra. Index-based object recognition in pictorial data management. *Computer Vision, Graphics, and Image Processing*, 52(3):416–436, 1990.

[4] Dengsheng Zhang and Guojun Lu. Review of shape representation and description techniques. *Pattern Recognition*, 37(1):1 – 19, 2004.

[5] Eamonn Keogh, Li Wei, Xiaopeng Xi, Michail Vlachos, Sang-Hee Lee, and Pavlos Protopapas. Supporting exact indexing of arbitrarily rotated shapes and periodic time series under euclidean and warping distance measures. *The VLDB JournalóThe International Journal on Very Large Data Bases*, 18(3):611–630, 2009.

[6] Thomas B Sebastian, Philip N Klein, and Benjamin B Kimia. Recognition of shapes by editing shock graphs. In *ICCV*, volume 1, pages 755–762, 2001.

[7] Masaki Hilaga, Yoshihisa Shinagawa, Taku Kohmura, and Tosiyasu L Kunii. Topology matching for fully automatic similarity estimation of 3d shapes. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 203–212. ACM, 2001.

[8] Faouzi Ghorbel, Stéphane Derrode, Sami Dhahbi, Rim Mezhoud, et al. Reconstructing with geometric moments. In *Proc. Int. Conf. on Machine Intelligence: ACIDCA-ICMI*, pages 5–7, 2005.

[9] PG Callahan, JP Simmons, and M De Graef. A quantitative description of the morphological aspects of materials structures suitable for quantitative comparisons of 3d microstructures. *Modelling and Simulation in Materials Science and Engineering*, 21(1):015003, 2012.

[10] Marcin Novotni and Reinhard Klein. Shape retrieval using 3d zernike descriptors. *Computer-Aided Design*, 36(11):1047–1062, 2004.

[11] Roger C Reed. *The superalloys: fundamentals and applications*. Cambridge university press, 2006.

[12] T Grosdidier, A Hazotte, and A Simon. Precipitation and dissolution processes in γ/γ' single crystal nickel-based superalloys. *Materials Science and Engineering: A*, 256(1–2):183 – 196, 1998.

[13] Judith Prewitt and Mortimer L Mendelsohn. The analysis of cell images. *Annals of the New York Academy of Sciences*, 128(1):1035–1053, 1966.

[14] Xiaofeng Ren and Jitendra Malik. Learning a classification model for segmentation. In *ICCV*, volume 1, pages 10–17, 2003.

[15] Christoph Sommer, Christoph Straehle, Ullrich Koethe, and Fred A Hamprecht. Ilastik: Interactive learning and segmentation toolkit. In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pages 230–233. IEEE, 2011.

[16] Verena Kaynig, Thomas Fuchs, and Joachim M Buhmann. Neuron geometry extraction by perceptual grouping in sstem images. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 2902–2909. IEEE, 2010.

[17] Ramakrishnan Mukundan and KR Ramakrishnan. *Moment functions in image analysisótheory and applications*. World Scientific, 1998.

[18] Adil Fahad, Najlaa Alshatri, Zahir Tari, Abdullah Alamri, Ibrahim Khalil, Albert Y Zomaya, Sebti Foufou, and Abdelaziz Bouras. A survey of clustering algorithms for big data: Taxonomy and empirical analysis. *IEEE transactions on emerging topics in computing*, 2(3):267–279, 2014.

[19] The hdbscan clustering library. http://hdbscan.readthedocs.io/en/latest/.

[20] Madeleine N. Kelly, Krzysztof Glowinski, Noel T. Nuhfer, and Gregory S. Rohrer. The five parameter grain boundary character distribution of -ti determined from three-dimensional orientation data. *Acta Materialia*, 111:22 – 30, 2016.

[21] McLean P Echlin, Naji S Husseini, John A Nees, and Tresa M Pollock. A new femtosecond laser-based tomography technique for multiphase materials. *Advanced Materials*, 23(20):2339–2342, 2011.

[22] D. G. Kendall. The diffusion of shape. *Advances in Applied Probability*, 9(3):428–430, 1977.

[23] David Marr and Herbert Keith Nishihara. Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London B: Biological Sciences*, 200(1140):269–294, 1978.

[24] Michael Brady. Criteria for representations of shape. *Human and machine vision*, 1, 1983.

[25] Robert J Woodham. Stable representation of shape. 1987.

[26] Thomas O Binford. Survey of model-based image analysis systems. *The International Journal of Robotics Research*, 1(1):18–64, 1982.

[27] Farzin Mokhtarian and Alan K Mackworth. A theory of multiscale, curvature-based shape representation for planar curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(8):789–805, 1992.

[28] Sven Loncaric. A survey of shape analysis techniques. *Pattern Recognition*, 31(8):983 – 1001, 1998.

[29] Herbert Freeman. Boundary encoding and processing. *Picture processing and psychopictorics*, 241, 1970.

[30] Li-De Wu. On the chain code of a line. *IEEE Transactions on pattern analysis and machine intelligence*, (3):347–353, 1982.

[31] Ernesto Bribiesca. A new chain code. *Pattern Recognition*, 32(2):235 – 251, 1999.

[32] Jukka Iivarinen and Ari JE Visa. Shape recognition of irregular objects. In *Photonics East'96*, pages 25–32. International Society for Optics and Photonics, 1996.

[33] Xiaolong Dai and Siamak Khorram. A feature-based image registration algorithm using improved chain-code representation combined with invariant moments. *IEEE Transactions on Geoscience and Remote Sensing*, 37(5):2351–2362, 1999.

[34] FREEMAN H Davis and H Freeman. A corner finding algorithm for chain coded curves. *IEEE Trans. On Computers*, 26:297–303, 1997.

[35] X.Y. Jiang and H. Bunke. Simple and fast computation of moments. *Pattern Recognition*, 24(8):801 – 806, 1991.

[36] Hong-Chih Liu and Mandyam D Srinath. Corner detection from chain-code. *Pattern Recognition*, 23(1-2):51–68, 1990.

[37] Debranjan Sarkar. A simple algorithm for detection of significant vertices for polygonal approximation of chain-coded curves. *Pattern Recognition Letters*, 14(12):959–964, 1993.

[38] Ernesto Bribiesca and Adolfo Guzman. Shape description and shape similarity measurement for two-dimensional regions. 1978.

[39] Babu M Mehtre, Mohan S Kankanhalli, and Wing Foon Lee. Shape measures for content based image retrieval: a comparison. *Information Processing & Management*, 33(3):319–337, 1997.

[40] William I Grosky, Peter Neo, and Rajiv Mehrotra. A pictorial index mechanism for model-based matching. *Data & Knowledge Engineering*, 8(4):309–327, 1992.

[41] Theodosios Pavlidis and Farhat Ali. Computer recognition of handwritten numerals by polygonal approximations. *IEEE Transactions on Systems, Man, and Cybernetics*, 5(6):610–614, 1975.

[42] Rajiv Mehrotra and William I Grosky. Shape matching utilizing indexed hypotheses generation and testing. *IEEE Transactions on Robotics and Automation*, 5(1):70–77, 1989.

[43] Vladimir I Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, pages 707–710, 1966.

[44] Raul S Montero and Ernesto Bribiesca. State of the art of compactness and circularity measures. In *International mathematical forum*, volume 4, pages 1305–1335, 2009.

[45] Daniel P. Huttenlocher, Gregory A. Klanderman, and William J Rucklidge. Comparing images using the hausdorff distance. *IEEE Transactions on pattern analysis and machine intelligence*, 15(9):850–863, 1993.

[46] William J Rucklidge. Locating objects using the hausdorff distance. In *Computer Vision, 1995. Proceedings., Fifth International Conference on*, pages 457–464. IEEE, 1995.

[47] Jingying Chen, Maylor K Leung, and Yongsheng Gao. Noisy logo recognition using line segment hausdorff distance. *Pattern recognition*, 36(4):943–955, 2003.

[48] M-P Dubuisson and Anil K Jain. A modified hausdorff distance for object matching. In *Pattern Recognition, 1994. Vol. 1-Conference A: Computer Vision & Image Processing., Proceedings of the 12th IAPR International Conference on*, volume 1, pages 566–568. IEEE, 1994.

[49] E Roy Davies. *Machine vision: theory, algorithms, practicalities*. Elsevier, 2004.

[50] Donald J Berndt and James Clifford. Using dynamic time warping to find patterns in time series. In *KDD workshop*, volume 10, pages 359–370. Seattle, WA, 1994.

[51] Gautam Das, Dimitrios Gunopulos, and Heikki Mannila. Finding similar time series. In *European Symposium on Principles of Data Mining and Knowledge Discovery*, pages 88–100. Springer, 1997.

[52] Tomasz Adamek and Noel O'Connor. Efficient contour-based shape representation and matching. In *Proceedings of the 5th ACM SIGMM international workshop on Multimedia information retrieval*, pages 138–143. ACM, 2003.

[53] Tomasz Adamek and Noel E O'Connor. A multiscale representation method for nonrigid shapes with a single closed contour. *IEEE Transactions on Circuits and Systems for Video Technology*, 14(5):742–753, 2004.

[54] Charles T Zahn and Ralph Z Roskies. Fourier descriptors for plane closed curves. *IEEE Transactions on computers*, 100(3):269–281, 1972.

[55] Gösta H Granlund. Fourier preprocessing for hand print character recognition. *IEEE transactions on computers*, 100(2):195–201, 1972.

[56] Eric Persoon and King-Sun Fu. Shape discrimination using fourier descriptors. *IEEE Transactions on systems, man, and cybernetics*, 7(3):170–179, 1977.

[57] Dengsheng Zhang and Guojun Lu. Shape-based image retrieval using generic fourier descriptor. *Signal Processing: Image Communication*, 17(10):825–848, 2002.

[58] CC Lin and Rama Chellappa. Classification of partial 2-d shapes using fourier descriptors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (5):686–690, 1987.

[59] Thomas R Crimmins. A complete set of fourier descriptors for two-dimensional shapes. *IEEE Transactions on Systems, Man, and Cybernetics*, 12(6):848–855, 1982.

[60] Yong Rui, Alfred C She, and Thomas S Huang. Modified fourier descriptors for shape representation-a practical approach. In *Proc of First International Workshop on Image Databases and Multi Media Search*, pages 22–23. Citeseer, 1996.

[61] Harry Blum. A transformation for extracting new descriptions of shape. In *Symposium on Models for the Perception of Speech and Visual Form*, 1967.

[62] Lorenzo Calabi and William E Hartnett. Shape recognition, prairie fires, convex deficiencies and skeletons. *The American Mathematical Monthly*, 75(4):335–342, 1968.

[63] Richard O Duda, Peter E Hart, et al. *Pattern classification and scene analysis*. Wiley New York, 1973.

[64] Kaleem Siddiqi, Ali Shokoufandeh, SJ Dickenson, and Steven W Zucker. Shock graphs and shape matching. In *Computer Vision, 1998. Sixth International Conference on*, pages 222–229. IEEE, 1998.

[65] Marcello Pelillo, Kaleem Siddiqi, and Steven W Zucker. Matching hierarchical structures using association graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(11):1105–1120, 1999.

[66] Benjamin B Kimia, Jackie Chan, Dale Bertrand, Seth Coe, Zachary Roadhouse, and Hueseyin Tek. Shock-based approach for indexing of image databases using shape. In *Voice, Video, and Data Communications*, pages 288–302. International Society for Optics and Photonics, 1997.

[67] Daniel Sharvit, Jacky Chan, Huseyin Tek, and Benjamin B Kimia. Symmetry-based indexing of image databases. In *Content-Based Access of Image and Video Libraries, 1998. Proceedings. IEEE Workshop on*, pages 56–62. IEEE, 1998.

[68] Guojun Lu and Atul Sajjanhar. Region-based shape representation and similarity measure suitable for content-based image retrieval. *Multimedia systems*, 7(2):165–174, 1999.

[69] Tom Huang, Sharad Mehrotra, and Kannan Ramchandran. Multimedia analysis and retrieval system (mars) project. *Digital Image Access & Retrieval [papers presented at the 1996 Clinic on Library Applications of Data Processing, March 24-26, 1996 Urbana-Champaign]*, 1997.

[70] Maytham Safar, Cyrus Shahabi, and Xiaoming Sun. Image retrieval by shape: a comparative study. In *Multimedia and Expo, 2000. ICME 2000. 2000 IEEE International Conference on*, volume 1, pages 141–144. IEEE, 2000.

[71] Kaushik Chakrabarti, Michael Ortega-Binderberger, Kriengkrai Porkaew, and Sharad Mehrotra. Similar shape retrieval in mars. In *Multimedia and Expo, 2000. ICME 2000. 2000 IEEE International Conference on*, volume 2, pages 709–712. IEEE, 2000.

[72] Johan WH Tangelder and Remco C Veltkamp. A survey of content based 3d shape retrieval methods. In *Shape Modeling Applications, 2004. Proceedings*, pages 145–156. IEEE, 2004.

[73] Ivan Sipiran, Benjamin Bustos, and Tobias Schreck. Data-aware 3d partitioning for generic shape retrieval. *Computers & Graphics*, 37(5):460–472, 2013.

[74] Ivan Sipiran and Benjamin Bustos. Harris 3d: a robust extension of the harris operator for interest point detection on 3d meshes. *The Visual Computer*, 27(11):963–976, 2011.

[75] Chris Harris and Mike Stephens. A combined corner and edge detector. In *Alvey vision conference*, volume 15, pages 10–5244. Citeseer, 1988.

[76] Alexei Elinson, Dana S Nau, and William C Regli. Feature-based similarity assessment of solid models. In *Proceedings of the fourth ACM symposium on Solid modeling and applications*, pages 297–310. ACM, 1997.

[77] Hari Sundar, Deborah Silver, Nikhil Gagvani, and Sven Dickinson. Skeleton based shape matching and retrieval. In *Shape Modeling International, 2003*, pages 130–139. IEEE, 2003.

[78] John Milnor. *Morse Theory.(AM-51)*, volume 51. Princeton university press, 2016.

[79] Georges Reeb. Sur les points singuliers díune forme de pfaff completement intégrable ou díune fonction numérique. *CR Acad. Sci. Paris*, 222(847-849):2, 1946.

[80] Dmitriy Bespalov, William C Regli, and Ali Shokoufandeh. Reeb graph based shape retrieval for cad. In *ASME 2003 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, pages 229–238. American Society of Mechanical Engineers, 2003.

[81] Yijun Xiao, Paul Siebert, and Naoufel Werghi. A discrete reeb graph approach for the segmentation of human body scans. In *3-D Digital Imaging and Modeling, 2003. 3DIM 2003. Proceedings. Fourth International Conference on*, pages 378–385. IEEE, 2003.

[82] Julien Tierny, Jean-Philippe Vandeborre, and Mohamed Daoudi. Partial 3d shape retrieval by reeb pattern unfolding. In *Computer Graphics Forum*, volume 28, pages 41–55. Wiley Online Library, 2009.

[83] Hermilo Sánchez-Cruz and Ernesto Bribiesca. A method of optimum transformation of 3d objects used as a measure of shape dissimilarity. *Image and Vision Computing*, 21(12):1027–1036, 2003.

[84] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Robotics-DL tentative*, pages 586–606. International Society for Optics and Photonics, 1992.

[85] Thomas Funkhouser, Patrick Min, Michael Kazhdan, Joyce Chen, Alex Halderman, David Dobkin, and David Jacobs. A search engine for 3d models. *ACM Transactions on Graphics (TOG)*, 22(1):83–105, 2003.

[86] SK Nayar and H Murase. Image spotting of 3d objects using parametric eigenspace representation. In *Proceedings of 9th Scandinavian conference on image analysis*, pages 325–332, 1995.

[87] Danny Roobaert and Marc M Van Hulle. View-based 3d object recognition with support vector machines. In *Neural Networks for Signal Processing IX, 1999. Proceedings of the 1999 IEEE Signal Processing Society Workshop.*, pages 77–84. IEEE, 1999.

[88] Ming-Kuei Hu. Visual pattern recognition by moment invariants. *Information Theory, IRE Transactions on*, 8(2):179–187, 1962.

[89] Venkatarama Krishnan. *Probability and random processes*. John Wiley & Sons, 2015.

[90] Nor Azura Md Ghani, Choong-Yeun Liong, and Abdul Aziz Jemain. Analysis of geometric moments as features for firearm identification. *Forensic science international*, 198(1):143–149, 2010.

[91] Masoud Alghoniemy and Ahmed H Tewfik. Geometric invariance in image watermarking. *IEEE transactions on image processing*, 13(2):145–153, 2004.

[92] Richard G Casey. Moment normalization of handprinted characters. *IBM Journal of Research and Development*, 14(5):548–557, 1970.

[93] K Udagawa, J Toriwaki, and K Sugino. Normalization and recognition of two-dimensional patterns with linear distortion by moments. *Electron. Commun. Japan*, 47(6):34–46, 1964.

[94] Mohamed Rizon, Yazid Haniza, Saad Puteh, Ali Yeon, Md Shakaff, Saad Abdul Rahman, Mamat Mohd Rozailan, Yaacob Sazali, Desa Hazri, and M Karthigayan. Object detection using geometric invariant moment. 2006.

[95] Alireza Khotanzad and Yaw Hua Hong. Invariant image recognition by zernike moments. *IEEE Transactions on pattern analysis and machine intelligence*, 12(5):489–497, 1990.

[96] Michael Reed Teague. Image analysis via the general theory of moments. *JOSA*, 70(8):920–930, 1980.

[97] Josef Bigun. Optimal orientation detection of linear symmetry, 1987.

[98] Sahibsingh A Dudani, Kenneth J Breeding, and Robert B McGhee. Aircraft identification by moment invariants. *IEEE transactions on computers*, 100(1):39–46, 1977.

[99] Jin Soo Noh and Kang Hyeon Rhee. Palmprint identification algorithm using hu invariant moments and otsu binarization. In *Computer and Information Science, 2005. Fourth Annual ACIS International Conference on*, pages 94–99. IEEE, 2005.

[100] Ji-Xiang Du, Xiao-Feng Wang, and Guo-Jun Zhang. Leaf shape based plant species recognition. *Applied mathematics and computation*, 185(2):883–893, 2007.

[101] Jan Flusser. On the independence of rotation moment invariants. *Pattern recognition*, 33(9):1405–1410, 2000.

[102] AG Mamistvalov. On the construction of affine invariants of n-dimensional patterns. *Bull Acad Science Georgian SSR*, 76(1):61–64, 1974.

[103] Alexander G Mamistvalov. N-dimensional moment invariants and conceptual mathematical theory of recognition n-dimensional solids. *IEEE Transactions on pattern analysis and machine intelligence*, 20(8):819–831, 1998.

[104] JP MacSleyne, JP Simmons, and M De Graef. On the use of moment invariants for the automated analysis of 3d particle shapes. *Modelling and Simulation in Materials Science and Engineering*, 16(4):045008, 2008.

[105] Jan Flusser and Tom Suk. Rotation moment invariants for recognition of symmetric objects. *IEEE Transactions on Image Processing*, 15(12):3784–3790, 2006.

[106] Tomas Suk and Jan Flusser. Graph method for generating affine moment invariants. In *Pattern Recognition, 2004. ICPR 2004. Proceedings of the 17th International Conference on*, volume 2, pages 192–195. IEEE, 2004.

[107] Tomáš Suk and Jan Flusser. Tensor method for constructing 3d moment invariants. In *Computer Analysis of Images and Patterns*, pages 212–219. Springer, 2011.

[108] C-H Teh and Roland T. Chin. On image analysis by the methods of moments. *IEEE Transactions on pattern analysis and machine intelligence*, 10(4):496–513, 1988.

[109] Yunlong Sheng and Lixin Shen. Orthogonal fourier–mellin moments for invariant pattern recognition. *JOSA A*, 11(6):1748–1757, 1994.

[110] N Canterakis. 3d zernike moments and zernike affine invariants for 3d image analysis and recognition. In *In 11th Scandinavian Conf. on Image Analysis*. Citeseer, 1999.

[111] Cyril Stanley Smith. The shape of things. *Scientific American*, 190(1):58–65, 1954.

[112] Cyril Stanley Smith. *Grain shapes and other metallurgical applications of topology*. American Society for Metals, 1951.

[113] Zachary Pincus and JA Theriot. Comparison of quantitative methods for cell-shape analysis. *Journal of microscopy*, 227(2):140–156, 2007.

[114] Guclu Ongun, Ugur Halici, Kemal Leblebicioglu, Volkan Atalay, Meral Beksaç, and Sinan Beksaç. Feature extraction and classification of blood cells for an automated differential blood count system. In *Neural Networks, 2001. Proceedings. IJCNN'01. International Joint Conference on*, volume 4, pages 2461–2466. IEEE, 2001.

[115] Amir Tahmasbi, Fatemeh Saki, and Shahriar B Shokouhi. Classification of benign and malignant masses based on zernike moments. *Computers in biology and medicine*, 41(8):726–735, 2011.

[116] Shubhi Sharma and Pritee Khanna. Computer-aided diagnosis of malignant mammograms using zernike moments and svm. *Journal of digital imaging*, 28(1):77–90, 2015.

[117] Khaled Khairy, JiJinn Foo, and Jonathon Howard. Shapes of red blood cells: Comparison of 3d confocal images with the bilayer-couple model. *Cellular and molecular bioengineering*, 1(2-3):173, 2008.

[118] Hunter Lee Elliott. *Cell Shape Invariant Quantification of Morphology, Migratino, and Subcellular Dynamics*. PhD thesis, Scripps Research Institute, April 2012.

[119] Francesco Merola, Pasquale Memmolo, Lisa Miccio, Roberto Savoia, Martina Mugnano, Angelo Fontana, Giuliana D'ippolito, Angela Sardo, Achille Iolascon, Antonella Gambale, et al. Tomographic flow cytometry by digital holography. *Light: Science & Applications*, 6(4):e16241, 2017.

[120] Elisabeth Truernit, Hélène Bauby, Bertrand Dubreucq, Olivier Grandjean, John Runions, Julien Barthélémy, and Jean-Christophe Palauqui. High-resolution whole-mount imaging of three-dimensional tissue organization and gene expression enables the study of phloem development and structure in arabidopsis. *The Plant Cell*, 20(6):1494–1503, 2008.

[121] Arabidopsis Genome Initiative et al. Analysis of the genome sequence of the flowering plant arabidopsis thaliana. *nature*, 408(6814):796, 2000.

[122] Richard Ernest Bellman. *Dynamic Programming*. Princeton University Press, 1957.

[123] Arthur Zimek, Erich Schubert, and Hans-Peter Kriegel. A survey on unsupervised outlier detection in high-dimensional numerical data. *Statistical Analysis and Data Mining*, 5(5):363–387, 2012.

[124] J. Ross Quinlan. Induction of decision trees. *Machine learning*, 1(1):81–106, 1986.

[125] Vladimir Vapnik. *Estimation of dependences based on empirical data*. Springer Science & Business Media, 2006.

[126] Simon Haykin and Neural Network. A comprehensive foundation. *Neural networks*, 2(2004):41, 2004.

[127] Sotiris B Kotsiantis, I Zaharakis, and P Pintelas. Supervised machine learning: A review of classification techniques. *Emerging artificial intelligence applications in computer engineering*, 160:3–24, 2007.

[128] Vladimir N Vapnik and Alexey J Chervonenkis. Theory of pattern recognition. 1974.

[129] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.

[130] Rich Caruana and Alexandru Niculescu-Mizil. An empirical comparison of supervised learning algorithms. In *Proceedings of the 23rd international conference on Machine learning*, pages 161–168. ACM, 2006.

[131] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.

[132] Geoffrey E Hinton and Sam T Roweis. Stochastic neighbor embedding. In *Advances in neural information processing systems*, pages 857–864, 2003.

[133] Martijn van Sebille, Laurens JP van der Maaten, Ling Xie, Karol Jarolimek, Rudi Santbergen, René ACMM van Swaaij, Klaus Leifer, and Miro Zeman. Nanocrystal size distribution analysis from transmission electron microscopy images. *Nanoscale*, 7(48):20593–20606, 2015.

[134] Walid M Abdelmoula, Karolina krkov, Benjamin Balluff, Ricardo J Carreira, Else A Tolner, Boudewijn PF Lelieveldt, Laurens van der Maaten, Hans Morreau, Arn MJM van den Maagdenberg, Ron MA Heeren, et al. Automatic generic registration of mass spectrometry imaging data to histology using nonlinear stochastic embedding. *Analytical chemistry*, 86(18):9204–9211, 2014.

[135] Otto Brinkkemper, Laurens van der Maaten, and Paul Boon. Identification of myosotis seeds by means of digital image analysis. *Vegetation history and archaeobotany*, 20(5):435, 2011.

[136] Laurens Van Der Maaten and Eric O Postma. Improving automatic writer identification. In *BNAIC*, pages 260–266, 2005.

[137] LJP Maaten and PJ Boon. Coin-o-matic: a fast and reliable system for coin classification. In *Proceedings of the MUSCLE Coin Workshop, Berlin, Germany*, pages 7–17, 2006.

[138] Ricardo JGB Campello, Davoud Moulavi, and Joerg Sander. Density-based clustering based on hierarchical density estimates. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 160–172. Springer, 2013.

[139] Stephen W Paddock. Confocal laser scanning microscopy. *Biotechniques*, 27:992–1007, 1999.

[140] Department of image processing: 3d rotation moment invariants. http://zoi.utia.cas.cz/3DRotationInvariants.

[141] Asa Ben-Hur and Jason Weston. A user's guide to support vector machines. In *Data mining techniques for the life sciences*, pages 223–239. Springer, 2010.

[142] Gudmund R Iversen and Helmut Norpoth. *Analysis of variance*. Number 1. Sage, 1987.

[143] Jacob Cohen. Multiple regression as a general data-analytic system. *Psychological Bulletin*, 70(6p1):426, 1968.

[144] Ivo Dinov. F-distribution tables. http://www.socr.ucla.edu/Applets.dir/F_Table.html.

[145] Richard Williams. Review of multiple regression. https://www3.nd.edu/~rwilliam/stats2/l02.pdf, Jan 2015.

[146] George AF Seber and Alan J Lee. *Linear regression analysis*, volume 329. John Wiley & Sons, 2012.

[147] Seongho Kim. ppcor: an r package for a fast calculation to semi-partial correlation coefficients. *Communications for statistical applications and methods*, 22(6):665, 2015.

[148] Myra L Samuels. Simpson's paradox and related phenomena. *Journal of the American Statistical Association*, 88(421):81–88, 1993.

[149] Herve Abdi. Multiple correlation coefficient. *The University of Texas at Dallas*, 2007.

[150] David L Bourell, Joseph J Beaman, Ming C Leu, and David W Rosen. A brief history of additive manufacturing and the 2009 roadmap for additive manufacturing: looking back and looking ahead. *Proceedings of RapidTech*, pages 24–25, 2009.

[151] Hideo Kodama. Automatic method for fabricating a three-dimensional plastic model with photo-hardening polymer. *Review of scientific instruments*, 52(11):1770–1773, 1981.

[152] JA Herbert. Solid object generation. *Journal of Applied Photographic Engineering*, 8(4):185–188, 1982.

[153] DL Bourell, JJ Beaman, HL Marcus, and JW Barlow. Solid freeform fabrication an advanced manufacturing approach. In *Proceedings of the SFF Symposium*, pages 1–7, 1990.

[154] Carl R Deckard. Method and apparatus for producing parts by selective sintering, September 5 1989. US Patent 4,863,538.

[155] Harris L Marcus, Joel W Barlow, Joseph J Beaman, and David L Bourell. From computer to component in 15 minutes: The integrated manufacture of three-dimensional objects. *Jom*, 42(4):8–10, 1990.

[156] JA Manriquez-Frayre and DL Bourell. Selective laser sintering of binary metallic powder. In *Proceedings of the Solid Freeform Fabrication Symposium, The University of Texas at Austin, Austin, TX*, volume 94, page 1990, 1990.

[157] JA Manriquez-Frayre and DL Bourell. Selective laser sintering of cu-pb/sn solder powders. *Marcus, H.; Beaman, J.; Bourell, D*, pages 252–260, 1991.

[158] G Zong, Y Wu, N Tran, I Lee, DL Bourell, JJ Beaman, and HL Marcus. Direct selective laser sintering of high temperature materials. In *Solid Freeform Fabrication Symposium Proceedings, The University of Texas at Austin*, pages 72–85, 1992.

[159] Gopalakrishna B Prabhu and David L Bourell. Supersolidus liquid phase selective laser sintering of prealloyed bronze powder. In *Proc. 4th Symp. on Solid freeform fabrication, Austin, TX, USA*, pages 317–324, 1993.

[160] Clyde O Brown, Edward M Breinan, and Bernard H Kear. Method for fabricating articles by sequential layer deposition, April 6 1982. US Patent 4,323,756.

[161] Peter Mercelis and Jean-Pierre Kruth. Residual stresses in selective laser sintering and selective laser melting. *Rapid prototyping journal*, 12(5):254–265, 2006.

[162] CR Knowles, TH Becker, and RB Tait. Residual stress measurements and structural integrity implications for selective laser melted ti-6al-4v. *South African Journal of Industrial Engineering*, 23(3):119–129, 2012.

[163] Amanda S Wu, Donald W Brown, Mukul Kumar, Gilbert F Gallegos, and Wayne E King. An experimental investigation into additive manufacturing-induced residual stresses in 316l stainless steel. *Metallurgical and Materials Transactions A*, 45(13):6260–6270, 2014.

[164] J-P Kruth, G Levy, F Klocke, and THC Childs. Consolidation phenomena in laser and powder-bed based layered manufacturing. *CIRP annals*, 56(2):730–759, 2007.

[165] William E Frazier. Metal additive manufacturing: a review. *Journal of Materials Engineering and Performance*, 23(6):1917–1928, 2014.

[166] EO Olakanmi, RF Cochrane, and KW Dalgarno. A review on selective laser sintering/melting (sls/slm) of aluminium alloy powders: Processing, microstructure, and properties. *Progress in Materials Science*, 74:401–477, 2015.

[167] Nikolay K Tolochko, Maxim K Arshinov, Andrey V Gusarov, Victor I Titov, Tahar Laoui, and Ludo Froyen. Mechanisms of selective laser sintering and heat transfer in ti powder. *Rapid prototyping journal*, 9(5):314–326, 2003.

[168] Mohd Rizal Alkahari, Tatsuaki Furumoto, Takashi Ueda, Akira Hosokawa, Ryutaro Tanaka, Abdul Aziz, and Mohd Sanusi. Thermal conductivity of metal powder and consolidated material fabricated via selective laser melting. In *Key Engineering Materials*, volume 523, pages 244–249. Trans Tech Publ, 2012.

[169] Hooyar Attar, Konda G Prashanth, Lai-Chang Zhang, Mariana Calin, Ilya V Okulov, Sergio Scudino, Chao Yang, and Jürgen Eckert. Effect of powder particle shape on the properties of in situ ti–tib composite materials produced by selective laser melting. *Journal of Materials Science & Technology*, 31(10):1001–1005, 2015.

[170] HJ Niu and ITH Chang. Selective laser sintering of gas and water atomized high speed steel powders. *Scripta Materialia*, 41(1):25–30, 1999.

[171] EO Olakanmi. Selective laser sintering/melting (sls/slm) of pure al, al-mg, and al-si powders: Effect of processing conditions and powder properties. *Journal of Materials Processing Technology*, 213(8):1387–1405, 2013.

[172] HP Tang, M Qian, N Liu, XZ Zhang, GY Yang, and J Wang. Effect of powder reuse times on additive manufacturing of ti-6al-4v by selective electron beam melting. *Jom*, 67(3):555–563, 2015.

[173] John A Slotwinski, Edward J Garboczi, Paul E Stutzman, Chiara F Ferraris, Stephanie S Watson, and Max A Peltz. Characterization of metal powders used for additive manufacturing. *Journal of research of the National Institute of Standards and Technology*, 119:460, 2014.

[174] ST Erdogan, PN Quiroga, DW Fowler, HA Saleh, Richard A Livingston, Edward J Garboczi, Peter M Ketcham, John G Hagedorn, and Steven G Satterfield. Three-dimensional shape analysis of coarse aggregates: New techniques for and preliminary results on several different coarse aggregates and reference rocks. *Cement and Concrete Research*, 36(9):1619–1627, 2006.

[175] John-Paul Latham, Antonio Munjiza, Xavier Garcia, Jiansheng Xiang, and Romain Guises. Three-dimensional particle shape acquisition and use of shape library for dem and fem/dem simulation. *Minerals Engineering*, 21(11):797–805, 2008.

[176] Edward J Garboczi. Three-dimensional mathematical analysis of particle shape using x-ray tomography and spherical harmonics: Application to aggregates used in concrete. *Cement and concrete research*, 32(10):1621–1638, 2002.

[177] ST Erdoğan, Aaron M Forster, Paul E Stutzman, and EJ Garboczi. Particle-based characterization of ottawa sand: Shape, size, mineralogy, and elastic moduli. *Cement and Concrete Composites*, 83:36–44, 2017.

[178] Michael A Taylor, Edward J Garboczi, ST Erdogan, and DW Fowler. Some properties of irregular 3-d particles. *Powder Technology*, 162(1):1–15, 2006.

[179] Lisa Markusson. Powder characterization for additive manufacturing processes, 2017.

[180] Jun Hao Tan, Wai Leong Eugene Wong, and Kenneth William Dalgarno. An overview of powder granulometry on feedstock and part performance in the selective laser melting process. *Additive Manufacturing*, 2017.

[181] Brian L DeCost, Harshvardhan Jain, Anthony D Rollett, and Elizabeth A Holm. Computer vision and machine learning for autonomous characterization of am powder feedstocks. *JOM*, 69(3):456–465, 2017.

[182] David G Lowe. Object recognition from local scale-invariant features. In *Computer vision, 1999. The proceedings of the seventh IEEE international conference on*, volume 2, pages 1150–1157. Ieee, 1999.

[183] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. In *Readings in computer vision*, pages 726–740. Elsevier, 1987.

[184] Gabriella Csurka, Christopher Dance, Lixin Fan, Jutta Willamowski, and Cédric Bray. Visual categorization with bags of keypoints. In *Workshop on statistical learning in computer vision, ECCV*, volume 1, pages 1–2. Prague, 2004.

[185] Martin Wattenberg, Fernanda Viégas, and Ian Johnson. How to use t-sne effectively. *Distill*, 1(10):e2, 2016.

[186] Chester T Sims. A history of superalloy metallurgy for superalloy metallurgists. *Superalloys 1984*, pages 399–419, 1984.

[187] RA Ricks, AJ Porter, and RC Ecob. The growth of γ' precipitates in nickel-base superalloys. *Acta Metallurgica*, 31(1):43–53, 1983.

[188] Yoshinao Mishima, Shouichi Ochiai, Noboru Hamao, Masayoshi Yodogawa, and Tomoo Suzuki. Solid solution hardening of nickel - role of transition metal and b-subgroup solutes. *Transactions of the Japan Institute of Metals*, 27(9):656–664, 1986.

[189] Xingming Zhang, Huiqiu Deng, Shifang Xiao, Xiaofan Li, and Wangyu Hu. Atomistic simulations of solid solution strengthening in ni-based superalloy. *Computational Materials Science*, 68:132 – 137, 2013.

[190] R. M. N. Pelloux and N. J. Grant. Solid solution and second phase strengthening of nickel alloys at high and low temperatures. *Transactions of the Metallurgical Society of AIME*, (218):232, 1960.

[191] W.L. Phillips. Mechanical properties of several nickel-platinum group metal alloys. *Transactions of the Metallurgical Society of AIME*, (230):526, 1964.

[192] D. Blavette, A. Bostel, and J. M. Sarrau. Atom-probe microanalysis of a nickel-base superalloy. *Metallurgical Transactions A*, 16(10):1703–1711, 1985.

[193] H. Harada, A. Ishida, Y. Murakami, H.K.D.H. Bhadeshia, and M. Yamazaki. Atom-probe microanalysis of a nickel-base single crystal superalloy. *Applied Surface Science*, 67(1–4):299 – 304, 1993.

[194] T Murakumo, Y Koizumi, K Kobayashi, and H Harada. Creep strength of Ni-base single-crystal superalloys on the γ/γ' tie-line. *Superalloys*, pages 155–62, 2004.

[195] J.X. Zhang, T. Murakumo, H. Harada, and Y. Koizumi. Dependence of creep strength on the interfacial dislocations in a fourth generation {SC} superalloy tms-138. *Scripta Materialia*, 48(3):287 – 293, 2003.

[196] Shouichi Ochial, Yoshihiro Oya, and Tomoo Suzuki. Alloying behaviour of Ni_3Al, Ni_3Ga, Ni_3Si and Ni_3Ge. *Acta Metallurgica*, 32(2):289 – 298, 1984.

[197] Minoru Doi, Toru Miyazaki, and Teruyuki Wakatsuki. The effect of elastic interaction energy on the morphology of γ precipitates in nickel-based alloys. *Materials Science and Engineering*, 67(2):247–253, 1984.

[198] B.H. Kear, A.F. Giamei, J.M. Silcock, and R.K. Ham. Slip and climb processes in γ' precipitation hardened nickel-base alloys. *Scripta Metallurgica*, 2(5):287 – 293, 1968.

[199] M Yamaguchi, V Paidar, DP Pope, and V Vitek. Dissociation and core structure of¡ 110¿ screw dislocations in l12 ordered alloys i. core structure in an unstressed crystal. *Philosophical Magazine A*, 45(5):867–882, 1982.

[200] R.W. Cahn, P.A. Siemers, J.E. Geiger, and P. Bardhan. The order-disorder transformation in ni3al and ni3alfe alloys—i. determination of the transition temperatures and their relation to ductility. *Acta Metallurgica*, 35(11):2737 – 2751, 1987.

[201] PH Thornton, RG Davies, and TL Johnston. The temperature dependence of the flow stress of the γ phase based upon ni 3 al. *Metallurgical and Materials Transactions B*, 1(1):207–218, 1970.

[202] TM Pollock and AS Argon. Creep resistance of cmsx-3 nickel base superalloy single crystals. *Acta Metallurgica et Materialia*, 40(1):1–30, 1992.

[203] V Sass, U Glatzel, and M Feller-Kniepmeier. Creep anisotropy in the monocrystalline nickel-base superalloy cmsx-4. *Superalloys*, pages 283–90, 1996.

[204] N Matan, DC Cox, P Carter, MA Rist, CMF Rae, and RC Reed. Creep of cmsx-4 superalloy single crystals: effects of misorientation and temperature. *Acta materialia*, 47(5):1549–1563, 1999.

[205] Frank RN Nabarro. Rafting in superalloys. *Metallurgical and Materials transactions A*, 27(3):513–530, 1996.

[206] Laura Jill Rowland. *Creep and microstructural stability of ruthenium-containing nickel-base single crystal superalloys*. 2005.

[207] Tresa M Pollock. *Creep deformation in nickel base superalloy single crystals.* PhD thesis, Massachusetts Institute of Technology, Department of Materials Science and Engineering, 1989.

[208] M. Fährmann, P. Fratzl, O. Paris, E. Fährmann, and William C. Johnson. Influence of coherency stress on microstructural evolution in model Ni-Al-Mo alloys. *Acta Metallurgica et Materialia*, 43(3):1007 – 1022, 1995.

[209] S.V. Prikhodko and A.J. Ardell. Coarsening of γ' in Ni-Al alloys aged under uniaxial compression: III. characterization of the morphology. *Acta Materialia*, 51(17):5021 – 5036, 2003.

[210] TM Pollock and AS Argon. Directional coarsening in nickel-base single crystals with high volume fractions of coherent precipitates. *Acta metallurgica et materialia*, 42(6):1859–1874, 1994.

[211] JS Van Sluytman and TM Pollock. Optimal precipitate shapes in nickel-base γ–γ alloys. *Acta Materialia*, 60(4):1771–1783, 2012.

[212] P. Caron and T. Khan. Improvement of creep strength in a nickel-base single-crystal superalloy by heat treatment. *Materials Science and Engineering*, 61(2):173 – 184, 1983.

[213] S Tin, TM Pollock, and WT King. Carbon additions and grain defect formation in high refractory nickel-base single crystal superalloys. *Ann Arbor*, 1001:48109, 2000.

[214] Q Feng, LJ Carroll, and TM Pollock. Soldification segregation in ruthenium-containing nickel-base superalloys. *Metallurgical and Materials Transactions A*, 37(6):1949–1962, 2006.

[215] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.

[216] Mary L Comer and Edward J Delp. Parameter estimation and segmentation of noisy or textured images using the em algorithm and mpm estimation. In *Image Processing, 1994. Proceedings. ICIP-94., IEEE International Conference*, volume 2, pages 650–654. IEEE, 1994.

[217] Michio Tokuyama, Kyozi Kawasaki, and Yoshihisa Enomoto. Kinetic equations for ostwald ripening. *Physica A: Statistical Mechanics and its Applications*, 134(2):323–338, 1986.

[218] Vassiliy A Epanechnikov. Non-parametric estimation of a multivariate probability density. *Theory of Probability & Its Applications*, 14(1):153–158, 1969.

[219] David R Bickel and Rudolf Frühwirth. On a fast, robust estimator of the mode: comparisons to other robust estimators with applications. *Computational Statistics & Data Analysis*, 50(12):3500–3530, 2006.

[220] Paul T Von Hippel. Mean, median, and skew: Correcting a textbook rule. *Journal of Statistics Education*, 13(2), 2005.

[221] Laura Jill Rowland. *Creep and microstructural stability of ruthenium-containing nickel-base single crystal superalloys.* 2005.

[222] JF Nie, ZL Liu, XM Liu, and Z Zhuang. Size effects of γ' precipitate on the creep properties of directionally solidified nickel-base super-alloys at middle temperature. *Computational Materials Science*, 46(2):400–406, 2009.