THE UNIFICATION OF STYLISTIC FORM & FUNCTION

by Ian Tseng

Presented March 2, 2011

Carnegie Mellon University

ABSTRACT

Great design is often the result of intelligent balancing of tradeoffs and leveraging of synergies between multiple product goals. While the engineering design community has numerous tools for managing the interface between functional goals in products, there are currently no formalized methods to concurrently manage stylistic form and functional requirements. The purpose of the work in this dissertation is to formalize ways to coordinate seemingly disparate but highly related goals of stylistic form and functional constraints in both computational design and for human designers. This work aims to provide a cohesive framework where both computational and cognitive findings are brought together to mutually inform and inspire the design process.

First, this problem was approached computationally with the development of an Artificial Neural Network based machine learning system that allows consumer judgments of stylistic form to be modeled quantitatively. Coupling this quantitative model of stylistic form with a Genetic Algorithm enables computers to concurrently account for multiple objectives in the domains of stylistic form and function within the same quantitative framework. This coupling then opens the door for computers to automatically generate products that not only work well, but also look good doing it.

Second, this problem was approached cognitively to explore ways to help human designers manage different goals of stylistic form and function more efficiently and effectively. An experiment was conducted which suggests that designers may sometimes have trouble fully utilizing knowledge they already have for managing different goals in design problems. This experiment shows evidence that analogical inspiration can help designers to overcome this knowledge block to more intelligently balance tradeoffs and leverage synergies in engineering design.

ACKNOWLEDGMENTS

First off I would like to thank my parents and sister for quite literally everything. Without them none of this could ever have been possible. They gave me morals and conviction, love and support, and every possible chance and opportunity that could be bestowed on me.

I would also like to thank my two esteemed advisors, Jonathan Cagan and Kenneth Kotovsky, for keeping it fun. Sure, their boundless knowledge, patient advice, genuine support, and zen-like guidance were no doubt invaluable to my journey. They had faith in me long after I did not, and shed doubt when my pie was in the sky. At the end of it all though, I think it was their fun banter that kept me here and kept me going.

Thanks to all of my lab mates, past and present, for the laughter and tears, the help and advise, the food and drinks, and for being guinea pigs through endless permutations of the same dull experiments. Between the lot of you, I estimate that you've judged upwards of 10,000 line drawings of hideous cars, and that's got to be worth something.

Special thanks to my committee for their tireless efforts, and endless patience.

A shout out to my best friend Drew, who helped me to waste my first two years here by having the time of my life.

Last, but certainly not least, I would like to thank Alison. You were always there to lend a shoulder to lean on, an ear to talk off, a brain to bounce ideas off, and a nosey to nom.

So long, and thanks for all the fish.

This work was supported by the National Science Foundation under grants: DMI0627894, BCS0717957, and SBE0738058.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER ONE

INTRODUCTION

1.1 Motivation

It is the pervading law of all things organic and inorganic,
Of all things physical and metaphysical,
Of all things human and all things super-human,
Of all true manifestations of the head,
Of the heart, of the soul,
That the life is recognizable in its expression,
That form ever follows function. This is the law.

This quote has been a cornerstone of modern architecture and design ever since Architect Louis Sullivan coined it in the late 19^{th} century. The idea is that form should never get in the way of the functional needs of a building or design. In engineering design the notions of aesthetic appearance and functional design are frequently seen as dichotomous rivals. Even at a corporate level, designing for the two tasks is frequently separated between two independent groups of people. Aesthetic form is usually handled by a group of industrial designers, while engineers typically handle the functional requirements. As a result, many products do not fully realize the potential of synergies or optimal tradeoffs between form and function. Much organizational research has shown that it is highly beneficial for these different groups of people who are working on different aspects of the same project to maintain an open dialogue and work together closely (Nadler et al., 1997; Jin & Levitt, 1996). Despite these organizational improvements, the complex interface between form and function is not fully understood. For one, how do these groups, or indeed even a single designer, interface between these two different goal types? Second, can a better understanding of how people solve design problems be used to better prepare or inspire designers to solve form and function design problems efficiently? Third, how much do form and function affect each other? Do people base their form judgments on perceived function? Fourth, how can computers co-represent form and function in a way that allows for mutual manipulation and quantitative optimization? This research makes inroads towards unifying these two goals in engineering design.

Applying Louis Sullivan's famous quote to design, the form of a design should be a result of functional needs, should look appropriate for its function, and should communicate its function to users. This relationship between form and function can be seen in action in the design of the traditional minivan. A minivan is not shaped the way it is for aesthetic reasons; rather the shape is a direct result of functional constraints and requirements, such as cargo and passenger capacity. Nevertheless, the designers of minivans are tasked with the goal of creating the most attractive shape that satisfies those goals.

This hierarchy of form following function has perhaps led to a problem in engineering design. Since functional constraints are often governed by theory and mathematical equations, they can generally be distilled into numerical target values or computer optimizable tradeoffs. As a result, it is often easy to see in the design process whether functional constraints have been met. Aesthetic form requirements on the other hand are not as simple. There are currently no quantifiable equations to reliably show how a particular design will look to the buyer, or automatically optimize stylistic design on a computer. This disparity can make it easier to allow for compromises in form rather than function in the design process. This problem is further compounded by the complexity of form preference. Not only is form preference difficult to quantify, but each person brings to the table a different desire of what the product should look like based on their needs, lifestyle, and goals for the product, as well as what society and culture have established as product norms and conventions. For example, the two digital cameras pictured in Figure 1.1 are both produced by Nikon and contain virtually the same photographic hardware inside their bodies [1], but they are marketed to two different target demographics. The camera on the left is geared towards laypeople who want a good simple camera. The casework is sleek, simple, modern, and easily pocketable. The camera on the right is designed for serious photographers (or people who want to feel like they are) who desire a smaller alternative to lugging around their professional cameras. The styling is utilitarian, classic, chunky, rugged looking, and shares many of the same style cues as Nikon's professional cameras. On online forums and reviews, both of these cameras have been lauded as being beautiful by their target demographics, yet they look substantially different.

Figure 1.1 – Two Nikon digital cameras targeted at average users (left) and prosumer users (right)

Product design has traditionally focused predominantly on the function of the product, with the form of the product often being approached as an afterthought after the technical specifications have been finalized (Nussbaum, 2002). The functional requirements of products continue to carry great importance, but in today's highly competitive market, a common key differentiator between functionally competent products is stylistic form. This is an even greater concern now that many products have increased in complexity, requiring the economies of scale to play a greater role, forcing companies to create multiple different products on a single shared platform. Volkswagen uses variations of a single platform to underpin nearly two dozen of their company's cars [2]. This platform sharing allows Volkswagen to cut development costs and create niche market vehicles that would normally be unable to sustain their own development costs. This highly beneficial arrangement would not be possible without a good working relationship between the stylistic form that differentiates the cars, and the functionality of the platform layout.

In a product designer's ideal world, it would be possible to modify and optimize all design parameters independently. Unfortunately, this is almost never the case. In reality, most design parameters are constrained by a variety of external concerns such as cost and technology, as well as interactive tradeoffs with other design parameters. It is often by properly balancing these interactive tradeoffs that truly great design emerges. For purely functional engineering specifications in design, methods like quality function deployment (QFD) exist which help to determine the goals of the product, how the competition meets those goals, importance of the goals to the consumer, and numerical targets to work toward (Ullman, 2003). Similarly, some work will be described in the background section that deals with the quantification and measurement of consumer preference towards pure stylistic form in products. Comparisons between stylistic form and function are currently viewed as comparing apples and oranges. Due to the amount of work that has been done studying the mutually exclusive goals of stylistic form and function, this work aims to couple these two goal types by focusing on the area of overlap between them, as shown in Figure 1.2. By better understanding how to manage the synergies and tradeoffs that exist between the two goals, new methods and models can be leveraged to better integrate stylistic form and function both cognitively and computationally in the design process, and change the hierarchy of form following function to a mutual relationship that balances their goals to optimally suit the needs and desires of the customer.

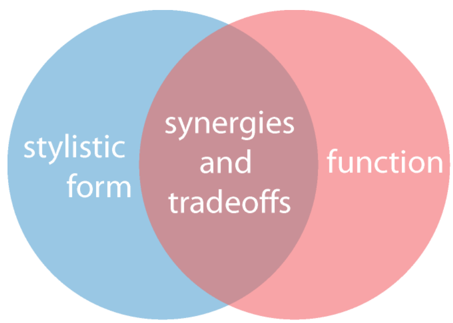

Figure 1.2 – The combination of stylistic form and function

1.2 THESIS

style *n. a distinctive appearance, typically determined by the principles according to which something is designed* (New Oxford American Dictionary, 2nd edition).

form *n.* the visible shape or configuration of something (New Oxford American Dictionary, 2nd edition).

function *n.* what a product or device must do (Wood & Greer, 2001).

The form and function of product shape are highly interrelated. When people are asked to make form based shape judgments, their reasoning stems from a wide range of considerations. A number of these considerations may be judged on implied functional performance based on the shape of the product. An example of this implied functionality is how people may perceive a car with thicker doors to be safer in collisions, regardless of whether the design is actually safer. The authors and contributors of a number of books (Parsons & Carlson, 2009; Ewen, 1990; Morgan, 1998; Harrison, 2001) have also reasoned that a source of form-based shape reasoning is based on stylistic norms. These stylistic norms include adherence to cultural and social fashions, and correlations with the media, society, and products they may have encountered before. Some of these correlations may have also originated from functional reasons and implications, but often have been largely confounded by stylistic norms. An example of the effect of stylistic norms is if a boxy car appears safer to a consumer because past designs have correlated boxy vehicle design with safety. In order to properly integrate these human product perceptions into the design process, these goals must be integrated with engineering considerations pertaining to the functionality and feasibility of the shape design.

We want to understand how designers can better use product shape to communicate to consumers, whether it is due to stylistic reasons or because it implies functional performance, and how this shape communication can be managed computationally with functional engineering requirements. It is important to understand both how designers approach product design problems, and how consumers react to changes in product shape. A better understanding of how people formulate and manipulate stylistic and functional judgments and the relationships between these judgments and functional engineering goals will lead to a better understanding of how functional performance goals can be met, and be better communicated to both consumers and designers. An improved understanding of the relationship between these judgments and goals shall allow potential synergies to be better exploited and tradeoffs to be better understood, revealing optimal tradeoff decisions. To follow the previous example, in addition to the importance of creating a safe vehicle design, it may also be important to communicate this feeling of safety to the consumer using functional form language. By understanding the relationship between creating a functionally safe vehicle design and one that communicates this feeling of safety, computers and designers will both be better equipped to manage any potential synergies and tradeoffs between the two goals more efficiently and effectively, resulting in a better design for the consumer.

The purpose of this work is to formalize ways to coordinate a variety of goals that are based on both stylistic form and function. The goal is not only to provide a synergy between stylistic form and function in design, but also to provide a cohesive framework where both computational tools and cognitive findings can be brought together to mutually inform and inspire the design process. The aforementioned goal results in the following thesis statement:

> *A better understanding of how people behave as both designers and consumers can inspire designers to better manage product design goals and enable computers to treat form and functional considerations in a unified fashion using the same quantitative methods.*

Three approaches are employed to examine how stylistic form and function can be unified in the design of products.

First, the relationships between different consumer judgments of stylistic form and function are studied and compared with actual measured functional performance for a better understanding of consumer style and performance perceptions and how accurately they correlate with measured performance. Understanding these relationships opens the door for better management of tradeoffs, potential synergies in the design process, and provides the foundation for the remaining two approaches.

Second, we examine this unification computationally. In order to better manage this tradeoff in computational optimization, it is necessary to develop software that allows a computer to automatically generate designs based on functional and stylistic form goals. This brings with it a number of fundamental challenges, not least of which is a current lack of methods for quantitatively treating stylistic form and human judgment. In order to accommodate this need, a method for computationally learning human judgments of stylistic form and function from survey data was developed using Artificial Neural Networks. This new ability to quantitatively model stylistic form and functional judgments, when coupled with Genetic Algorithms, enables computers to generate stylish and functionally communicative designs without further human interaction. This unification of stylistic form and function also enables computers to couple stylistic form judgments with functional targets to automatically generate designs that satisfy goals from both goal categories.

Third, this unification is examined cognitively by studying ways to enrich and inspire human designers to solve form and function design problems more efficiently and effectively. An experiment was conducted that studied the effect of how differing representations of inspirational knowledge affect the fluency with which human designers mentally manage multiple stylistic and functional goals in a design problem. Findings from this experiment led to insights about how to better prepare designers with applicable problem representations for more efficient management of conflicting goals and tradeoffs in a form and function design problem.

The knowledge gained from these three approaches culminates to help suggest new methods, potential improvements, and suggestions for future work for both computational optimization and human designers towards solving stylistic form and functional problems.

This work does not make any attempt to create absolute definitions of aesthetics, style, form or beauty; rather this work examines and leverages the underlying relationships that connect consumer judgments of stylistic form in products with quantifiable attributes of product form. The experiments in this dissertation require a simple design problem with a richly interactive relationship between stylistic form and function. This design problem must be applicable and familiar to all participants. A vehicle styling and functional shaping problem is chosen as an illustrative platform for testing this interface between form and function, although the methods are kept general and amenable to a wide variety of applications and product types.

1.3 DISSERTATION OUTLINE

The three approaches discussed in the previous section and their supporting backgrounds have been divided into the following chapters.

Chapter 2 reviews relevant research to the goal of coupling stylistic form and function in design. This literature review is split up into two sections that focus on the computational and cognitive approaches taken in this dissertation.

Chapter 3 presents the computational models that are used throughout this dissertation as an illustrative example for the work discussed and to generate the images used for stimuli in the experiments conducted. The models presented in this chapter are not the focus of the contributions in this dissertation; rather they are a necessary means to deliver those contributions. The work discussed in this dissertation can be applied to any product type that can be represented using the methods discussed in this chapter.

Chapter 4 focuses on better understanding consumers and their judgments about products for use as a foundation to the work done in Chapters 5 and 6. Towards this goal, an experiment was conducted to study the relationships between different consumer judgments of stylistic form and function, that compared them with actual measured functional performance for a better understanding of consumer style and performance perceptions and how accurately they correlate with measured performance.

Chapter 5 presents a series of three methodologies and associated experiments that quantitatively model consumer judgments. This quantification allows these consumer judgments to be treated using the same methods as functional goals in computational optimization. This chapter concludes with a method that leverages Artificial Neural Networks and Genetic Algorithms to computationally model and generate vehicle designs that satisfy both stylistic form goals and functional goals.

Chapter 6 describes two experiments that were conducted to study the circumstances that best allow designers to successfully apply and leverage their knowledge to solve design problems efficiently and effectively.

Chapter 7 summarizes the work in this dissertation, discusses future directions that this research should explore, and the contributions made towards the body of knowledge that informs design.

CHAPTER TWO
BACKGROUND

2.1 COMPUTATIONAL BACKGROUND

A vast amount of work has been done examining, measuring, and trying to model consumer preference computationally. One approach to capturing the effect of consumer preference is seen in a decision-based design framework proposed by Hazelrigg (1998). Hazelrigg's model assumes that engineering design is a decision-making process, and can be simplified into two steps: determine all possible design options, and choose the best one. This simplification allows many well-developed theories and methods from economics, operations research, decision sciences, and other disciplines to be applied to engineering design. Many researchers have since extended and built off this framework to evaluate customer preferences for design alternatives. Li and Azarm (2000) developed a system that models survey data using conjoint analyses (Luce & Tukey, 1964; Green & Rao, 1971; Johnson, 1974; Green & Srinivasan, 1978) to create utility functions that explicitly measure consumer preferences and account for them in the design optimization of cordless screwdrivers. Wassenaar and Chen (Wassenaar & Chen 2003, Wassenaar et al. 2005) also extended Hazelrigg's framework and used discrete choice analysis with revealed preference data to create a logit model to predict expected profit as a function of product attributes and demographic information. Michalek et al. (2004, 2006) applied the method of analytical target cascading decomposition to coordinate models of engineering and market performance. This method used a mixed logit model to determine consumer preference of bathroom scale designs. These consumer preferences were then balanced with a cost model to optimize the design for maximum profit, taking market considerations into account. While this bathroom scale model did account for the dimensions of the scale, as well as the size and formats of the dial readout, none of these methods truly consider the stylistic form of the product, or the consumer's aesthetic preference.

In order to incorporate stylistic form, a method must be employed for generating shape in a controlled deterministic manner that allows for a connection between the shape being generated and the measured preference. One such possible computational design tool is the use of shape grammars (Stiny 1980). A shape grammar is a set of shape transformation rules that allows for an initial shape to be altered in constrained ways to generate a desired shape. Over the years shape grammars have been developed to generate designs ranging from architectural floorplans (Stiny & Mitchell, 1978) and row houses (Cagdas, 1996), to coffeemakers (Agarwal & Cagan, 1998). Shape grammars have also been used to define branding cues of Harley Davidson motorcycles (Pugliese & Cagan, 2002) and Buick cars (McCormack et al., 2004). Orsborn (Orsborn et al., 2006; Orsborn, 2009a; Orsborn, 2009b; Orsborn et al., 2010) also built a shape grammar for designing cars, which he pared down to create a simpler seven parameter model to design the front end of SUVs. This model was used to run a discrete choice consumer survey to gauge aesthetic preference. This survey data was used to generate utility functions on each of the seven parametric attributes. Using these utility functions, new SUV front ends could be generated that should give high, medium, and low utility to each specific survey-taker. In a verification survey, it was found that the high utility option was selected by survey-takers 78% of the time, and the low utility option was only selected 2.3% of the time, showing that the consumers' aesthetic preference was modeled faithfully. One potential drawback to inferring form preference using utility functions is an inherent difficulty at modeling interaction effects between parameters (Brambor et al, 2005; Green, 1984; Hagerty, 1986), which is crucial in the study of aesthetic preference (Green & Srinivasan, 1990). Since human judgment of aesthetic form is complex and often formed with a combination of design elements, the ability to more easily model interaction effects with greater resolution in highly modal data would be invaluable, and is a focus of the current work.

The literature up until this point has predominantly examined the blanket issue of preference, but in reality preference is made up of a number of lower-level judgments and feelings. The decomposition and separate treatment of these reasons for why consumers prefer a product will allow for more powerful and universal grasp on preference as a whole. One reason for this is that preference differs dramatically between different consumers, while judgments of more specific product traits, such as whether a product appears to be modern or rugged, are generally less varied. Because of this, it may be useful to model these more specific judgments separately first, so that they can be recombined later to target specific combinations of consumer needs and desires. For instance, if a customer would prefer that a product appear rugged and not modern, their preference can be modeled as a composite of those two more specific judgments.

One method that can be used to decompose product goals is Kansei engineering (Nagamachi, 2002; Bouchard et al., 2003). Kansei engineering is a tool developed by Mitsui Nagamachi in the 1970s where the desired emotional feelings that result from interacting with the product can be extracted, decomposed, and used to define tangible product functions and design goals. The Japanese word "Kansei" roughly translates as "emotional feeling". Kansei engineering was first used in industry for the development of the Mazda Miata roadster. The zero-level concept, or overall target emotion for the Miata was defined as "unification of driver-machine". This zero-level concept was then decomposed into 1st-level to nth level concepts of more specific feelings, such as "speedy feeling". The decomposition continues until the desired concepts can be assigned required physical traits, such as engine noise frequency spectrums. These required physical traits are often measured using surveys that employ the semantic differential. The semantic differential (Osgood et al., 1957) is a way to measure the meaning of concepts and break them down into multiple dimensions of opposing adjectives using factor analysis. Due to the semantic differential's capability with abstract concepts, a common recent use is on surveys where consumers are asked to rate a product on a linear scale that ranges between two opposing words. For instance a variety of engine noises can be played, and the participant can be asked to score the noises in a linear range between exciting and bland, powerful and weak, etc.

In the examination of brand by Chen et al. (2004), six commonly occurring adjectives were selected from shampoo advertisements to represent common brand goals for shampoo. They then measured participants' ratings using the semantic differential on how well these adjectives described eighteen different hair care product bottles, which were stripped of all labels and all painted the same color. Using the resulting survey data, a shape grammar was developed describing the different needs of shampoo bottle shaping and targeted brand goals.

The semantic differential was also used by Achiche and Ahmed (2008) to model how a shape would be perceived by participants. Based on simple geometric relationships and survey data, a fuzzy logic system was implemented that could predict whether eleven of the twelve shapes appeared aggressive or soft. One criticism of this research is that the shapes used were designed specifically to be either extremely aggressive or soft with no middle ground. Other research that examine survey concerns in design are Chang et al. (2003) who tracked the differences between participants of different age groups on their product form perceptions, and Yanagisawa and Murakami (2008) who studied the effects of viewpoint shifts and latent sensitivity of customers when evaluating shape aesthetics.

Reid et al. (2009) presented a methodology for quantifying the perceived environmental friendliness of vehicle silhouettes. In this paper, participants rated computer generated vehicle designs on how environmentally friendly the design appeared, or how likely the designs were deemed to be inspired by nature. This data was correlated with the actual physical positions used in each vehicle design. This work found that vehicles with shape discontinuities, leading to a boxier shape, were less likely to be perceived as inspired by nature, and that vehicles with more raked front and rear windscreens, and more gradual transitions into the roofline, in turn are more likely to be perceived as environmentally friendly.

Morphing is a less computationally intense alternative to shape grammars for shape generation (Hsiao & Liu, 2002; Smith et al., 2007). Morphing starts with a population of known designs that are parametrically created, and by taking linear averages between the parameter values between two or more designs, a new linear recombination of the parent designs can be created. Since all parameter values lie within the range of the parent designs, the system is stable, meaning that all possible morphed combinations will be geometrically valid. A downside to this stability is that designs outside of the realm of the parent population cannot be created. Hsiao and Liu (2002) combined morphing and the semantic differential based surveys to create and judge LCD monitor designs. Using a grey theory prediction model, the ratings of previously unseen monitor designs could be determined computationally.

Smith et al. (2007) used morphing methods and an original population of 30 production cars to explore the design space of car shapes. Using this morphing model, it was possible for them to approximate the curb weight of vehicles based on known curb weights of the original population. A number of designs were self-scored for “sportiness”, which allowed a transformation vector to be derived that could be added to the parameter values of any design to make that vehicle appear more sporty. This method of adding a transformation vector allows these morphed models to leave the safe boundaries of the original population of car designs, but also opened the door to possibly invalid solutions and model instability.

Another interesting method for generating form was attempted by Yannou et al. (2008), who took a unique approach towards the requirements of shape generation by using a Fourier series to represent the genetics of car silhouettes. The car designs could then be refined to better fulfill user desires using an Interactive Genetic Algorithm (IGA). A Genetic Algorithm (GA) is a genetically inspired optimization technique that creates a random population of designs, and refines the population for improved performance by calculating the fitness of each design to probabilistically remove weaker design attributes and reinforce stronger design attributes. IGAs, also called Interactive Evolutionary Computation (IEC), are genetic algorithms that specifically employ the use of a human to perform the fitness judgments. IGAs are a very powerful method to allow people with little or no CAD experience to interact with software to create and modify designs and shapes. IGAs are particularly well suited in the generation of stylistic form and have been used successfully in the fields of fashion design (Kim & Cho, 2000), and in the shape design of eyeglass frames (Yanagisawa & Fukuda, 2004) A survey paper that covers many applications and concerns for IGAs can be found in (Takagi, 2001). A major limitation of IGAs is that a human is required to perform the fitness evaluation, which violates two large goals of computational design – to reduce human involvement and to speed up processing time.

Thanks to advances in artificial intelligence, it is possible to train a computer to closely mimic human responses to stylistic form. An Artificial Neural Network (ANN) is a computational model that contains a network of interconnected artificial neurons or nodes arranged in layers (Mitchell, 1997). This method was inspired and is loosely based on how biological neural networks function. ANNs have been shown to have excellent abilities to recognize patterns in large data sets (Fukushima, 1980, 1984) and have been used successfully to perform character recognition (Fukushima, 2003). ANNs have been coupled with a parametric model and the semantic differential method to model human preference in office chair design (Hsiao &

Huang, 2002) and bridge design (Yasuda et al., 1995). Hsiao and Huang trained an ANN to match a combination of discrete and continuous parametric model values to continuous survey responses. After training, the ANN could accurately mimic the survey takers and rate the chair designs that were not in the training data set on the six stylistic form judgments asked on the survey. Yasuda et al used five-level aesthetic surveys to judge 140 photographs of bridges on assessments of how impressive, unique, sophisticated, friendly, harmonious, and beautiful each bridge looked to survey takers. Structural information about each surveyed bridge design was also recorded. The resulting survey data was used to train an ANN to predict how people would perceive the aesthetics of a bridge based purely on structural data. The results of this study were unfortunately inconclusive, which the researchers attributed to vagueness of the aesthetic terms and inconsistency between survey takers. One limitation when trying to generate new designs using ANNs is their unidirectionality, meaning that while the neural network can mimic the decisions in the training data, it is impossible to reverse the direction and decode the input parameters from a desired output state. This differs from the utility function methods described previously, which can be used directly to generate new designs. What is needed to complete the picture with ANNs is a method for network inversion.

This desired network inversion can be achieved by coupling an artificial neural network with a genetic algorithm. An ANN can be used as the fitness function of the GA, which negates the need for the human fitness judgments as used in IGAs. This results in improved autonomy and speed, allowing new designs to be automatically bred in the GA that satisfy a desired output of the ANN. This method was used by Baluja et al (1994) to evolve and generate computer images that were pleasing to users. This method has also been used with limited success on generating musically pleasing phrases in jazz music (Griffith & Todd, 1999; Biles et al., 1996). Bull (1999) further refined the technique into a model-based evolutionary computation system that uses a generic parametric model form, and is iteratively retrained using the explicit fitness function every fixed number of generations to retrain the neural network for greater accuracy around the optimal solution. It is assumed in this system that the fitness function is costly, meaning that it is either time or resource intensive to evaluate, so minimizing how often it needs to be evaluated is of large benefit. This is shown to be much more efficient and effective than other methods that retrain on other schedules, or that do not retrain at all. Coupling ANNs and GAs with survey data built from five-level surveys to learn aesthetic form, much like the work done here has been attempted by Tsutsumi and Sasaki (2008) who paired this method with a five-level Kansei evaluation survey to design a gymnasium roof structure that satisfied both goals of aesthetic beauty and calculations of maximum stress. The computational work in this dissertation will build on the methods used and developed by Bull, and Tsutsumi and Sasaki.

Due to the wide availability of relevant literature that deals with functional optimization, this literature review has predominantly focused on the representation, surveying, modeling, and generation of stylistic form. A number of the methods described above can also be used to model and generate designs based on functional requirements. Since structural strength and aerodynamics are two functional requirements that interface greatly with the shape and stylistic form of a design, elegant methods of modeling or generating based on these needs are also of great interest. GAs have been used extensively to evolve designs based on functional requirements such as structural optimization (Adeli & Kumar, 1995) and for generating optimal aerodynamic shapes (Olhofer et al., 2001). Since computers can be very accurate at calculating structural strength or performing

Computational Fluid Dynamics (CFD) analyses, these computer driven finite element fitness functions are well suited for most functional optimization. Even still, calculations for CFD analyses are costly and time consuming, especially for large and complicated systems where computation time can take upwards of several months (Becker et al. 2007). A similar coupling of ANNs and GAs to what was discussed earlier was applied to CFD analyses by Becker et al. (2007). Using this method a multilayer ANN was trained on data from 72 CFD simulations and was used as the fitness function for a GA that could then generate new shapes that satisfy the desired result of the CFD analyses to within 10%. This type of system can prove useful when a quick CFD approximation is needed, such as in early design stages, or to give a user real-time feedback. A method similar to this has also been used in the field of fire safety (Yuen et al., 2004; Lee et al., 2004).

A wide range of methods have been discussed for generating and visualizing stylistic form, surveying and learning human preferences and judgments, and generating new designs. In order to tie these methods and models together and create a unified system that can cooperatively treat form and function, multiple fitness functions of various formats will have to be balanced and treated together. A number of relevant papers that cover multi-objective genetic algorithm techniques and their implementation include Deb et al. (2002), Coello (2000), Fonseca and Fleming (1995), Grignon and Fadel (2004), Gantovnik et al. (2006), and Zitzler and Thiele (1999).

2.2 Cognitive Background

The notion of style is richly pervasive throughout human history, and has had a wide range of implications, contexts, and uses throughout that history. The word style has been used to characterize trends in everything from ancient Greek writing to film noir movies, but it is a term that has eluded hard definition (Eckert & Do, 2006). In a review of what psychology and computational research can tell us about style, Stacey (2006) discusses how the style perception of an object involves the interaction of perceptual and conceptual processes to understand the object. He goes on to state that this understanding is based on combining visual properties with structure, function, and behavior. In order to properly inform the design process, two dichotomous understandings of style need to be balanced. The first is a look from the outside of how the object looks. The second is a look from the inside at how an object is made, and how it accomplishes its purpose. From this research it can be extended that style is dependant on both form and function, and that form and function appear to represent two separate but related representations or problem spaces that will need to be considered together.

Maher and Tang (2003) also assume two parallel search spaces in a cognitive model of design, but they parse the space into the problem space and the solution space. Two protocol studies were conducted wherein the subjects were asked to design an electric kettle. Monitoring the transcripts showed evolution of the problem definition, the solution definition, and switches back and forth between the two search spaces. A computational model was then proposed to operate iteratively with a genetic algorithm by refining each search space by using the other as a fitness function to evaluate alternative designs. These two search spaces eventually converge with a final solution.

Klahr and Dunbar (1988) similarly view scientific reasoning as being characterized by two problem spaces, a model they call Scientific Discovery as Dual Search (SDDS). SDDS breaks problem spaces down into two spaces, the hypothesis space, and an experiment space. In the hypothesis space, participants search their memory and form hypotheses based on prior knowledge. In the experiment space, participants form hypotheses based on the results of previous experiments. The model also describes how the problem-solver switches between the two problem spaces and how information gained from each mutually interface.

Dunbar (1993) later found that having different sub-goals in the same problem can constrain cognitive processes and have a profound effect on the way people mentally represent and solve problems. Participants in the experiment were taught some basic facts about microbiology, and were asked to discover how genes control other genes in a computer simulation. The participants were then given one of two possible goals. One group was asked to search for evidence consistent with the initial hypothesis, while the other group searched for evidence discrepant with the initial hypothesis. He hypothesizes that the two different sub-goals frame the problem differently and thus create different mental representations of the problem. A subset of participants in the group that focused on discrepant evidence were able to solve the problem, while none of the participants in the group that focused on consistent evidence did.

While these mental models do not explicitly represent form and function, these experiments suggest that people do in fact create and maintain multiple representations of problems they are trying to solve. This knowledge can aid in understanding how designers create mental models and switch between multiple goals of form and function. Another important factor to consider is how the schedule of switching between the two representations might affect creativity or time to solution. In a study of task rotation, Madjar and Oldham (2006) examined the effect of having participants either cycle in four fixed 3 minute intervals through three idea generation tasks, or worked serially for 12 minutes on each of the three tasks. By monitoring how many tasks each participant preferred to be engaged in at a time, it was found that participants fell into two categories. The group of participants who preferred working on multiple tasks at the same time performed better when they were given the opportunity to rotate tasks, while the group of participants who preferred working on fewer tasks at a time performed better when they were to finish one problem before starting the next.

Research has theorized (Amabile, 1983; Crosby, 1968; Guetzkow, 1965) or found (Shalley, 1991) that giving participants the freedom of high personal discretion on which of multiple problems they work on at a given time increases their productivity and the creativity of their solutions. Extending the findings of this experiment, Madjar and Shalley (2008) conducted an experiment to investigate the effects of working on multiple tasks with multiple goals at the same time. Participants in the first condition were given the freedom to switch between tasks at their own discretion, while the computer limited the total time on each problem to ten minutes. Participants in the second condition were given the three problems serially, with ten minutes available to solve each problem. It was hypothesized that the higher creativity that resulted from giving the participants discretion to switch tasks was a result of forced incubation, of having active goals for all three tasks.

But before looking at how participants represent multiple goals, it is important to gain insight on how people solve problems with a single goal, and how they bring new information into problem solving. Similarly, even when solving a problem with a single goal, there are multiple ways to mentally represent the problem. Research has shown that showing a problem solver an example solution can fixate the problem solver into a single representation of the problem and limit creativity (Chrysikou & Weisberg, 2005; Jansson & Smith, 1991; Perttula & Liikkanen, 2006; Purcell & Gero, 1996). Designers have even been observed to incorporate poor aspects of existing solutions into their own solution (Jansson & Smith, 1991). These previous experiments show evidence for fixation caused by pictorial examples given with the design problem.

People are more likely to incorporate relevant information when they encounter the information after they have started the problem. In the case where the problem-solver has the goal to solve a problem but has not yet completed the solution, the problem solver has an open problem-solving goal. Having an open goal actually makes it more likely that relevant information is incorporated into problem solving even when the person is not actively engaged in solving the problem (Moss et al., 2007, 2008). Moss et al. defined an open goal as a goal which has been set but one for which the associated task has not been completed (i.e., work on the task has been interrupted or suspended). They found that implicit hints were more likely to affect problem solving when presented after an initial unsuccessful attempt at the problem as compared to presenting the hint before the initial attempt at the problem, which suggests that open goals can make it easier for participants to assimilate and incorporate hints into problem solving. In order to study the effects of open goals and fixation on design problem solving, an experiment was conducted and is described Section 6.1.

CHAPTER THREE
COMPUTATIONAL REPRESENTATIONS OF FORM AND FUNCTION

The computational representations and evaluations described in this chapter are used throughout the remainder of this dissertation, although sometimes with minor modifications as noted. These representative models are not the focus of the contributions in this dissertation; rather they are a necessary means to deliver those contributions. The work presented in this dissertation can be applied to the design of a wide variety of product categories. Nevertheless, in order to elicit responses from consumers and to better illustrate the abilities of the methods, an illustrative design problem is needed. An appropriate design problem should be one with a richly interactive relationship between stylistic form and function, and should also be applicable and familiar to all participants. A vehicle styling design problem will be used as an illustrative platform for testing the relationship between stylistic form and function, although the methods will be kept general and amenable to a wide variety of applications.

A common thread that runs all phases of this work is the desire to bridge the gap between stylistic form and functional goals. There is an inherent disconnect between these two goal types due to differences in representation. Form is represented graphically, whereas functional goals are usually represented numerically. In order to incorporate both goal types, a method must be employed for generating shape in a controlled deterministic manner that connects the shape being generated with a unique numerical representation, which enables direct and objective comparisons with functional goals, as well as the ability for shape representation and computer algorithms to interface. As discussed in the literature review, there are many methods to accomplish this goal, including shape grammars (Stiny 1980; Cagdas, 1996; Agarwal & Cagan 1998; Pugliese & Cagan, 2002; Orsborn et al., 2006; Orsborn, 2007), morphing (Hsiao & Liu, 2002; Smith et al., 2007), and parametric models (Hsiao & Huang, 2002). A parametric model for generating cars is described in Section 3.1. In addition to being able to generate car designs, a method is also needed to assess the functional performance of each car design. An aerodynamic model, a volumetric estimator, and a center of gravity estimator were developed and are described in Sections 3.2 and 3.3.

3.1 PARAMETRIC DESIGN GENERATION

In order to consistently generate a large variety of car designs, a parametric car design model was developed in both MATLAB and C++. This model consists of eight cubic Bézier curves and two circles as shown in Figure 3.1. The overall length of the vehicle is constrained, so all parametric changes only affect the height and proportions of the vehicle, and not the length. A Bézier curve is a parametrically defined curve that is defined by control points and is commonly used in computer graphics. One useful characteristic of Bezier curves is that the control points define the end points of the curve, the direction of the curve at the ends, and the magnitude of the influence of that magnitude. In this model, cubic Bézier curves are used, which are defined by four control points, two of which are the endpoints. Curve 1 defines the front bumper, Curve 2 defines the hood, Curve 3 defines the front windshield, Curve 4 defines the roof, Curve 5 defines the rear windshield, Curve 6 defines the trunk lid, Curve 7 defines the rear bumper, and Curve 8 defines the floor pan of the car. These eight curves are defined by twelve design parameters that can be varied continuously between 0 and 100. These design parameters are represented as a vector of twelve integers, which will hereby be called the *chromosome*. The vehicle pictured in Figure 3.1 represents a chromosome with all parameters at a value of 50. The chromosome for this vehicle is shown below the vehicle design in bold, with the twelve parameters labeled as P1 – P12. The design parameters are discussed in Section 3.1.1, and the relationship between the design parameters and the eight curves are discussed in Section 3.1.2. Sample designs created with this model are shown in Section 3.1.3.

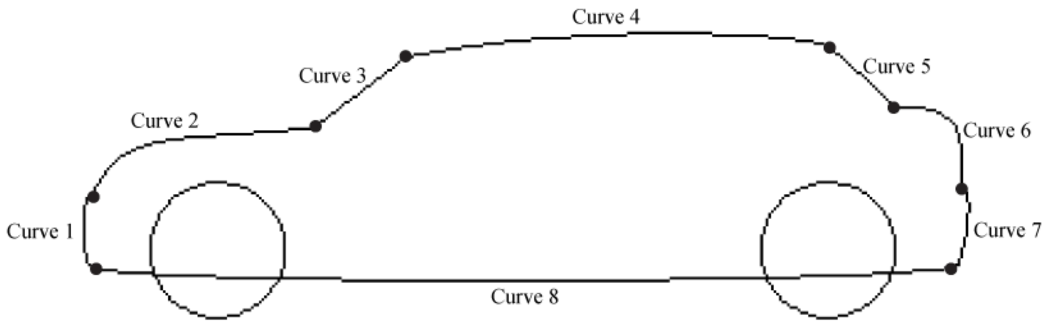

	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	
[	**50**	**50**	**50**	**50**	**50**	**50**	**50**	**50**	**50**	**50**	**50**	**50**	]

Figure 3.1 – Numbering of Bezier Curves used to draw car profiles, and the chromosome associated with the car design

3.1.1 DESIGN PARAMETERS

The twelve design parameters that make up a car design's chromosome in this parametric vehicle design model are described in this section. Each parameter can vary from 0 – 100, and help to define certain characteristics of the overall vehicle shape. The four control points of each of the eight curves that make up vehicle model are defined based on the value of parameters and the position of other curves. The chromosome for each of the illustrating designs pictured in this section is a modification of the chromosome in Figure 3.1, except for the one parameter being illustrated, which has a value of 0 for the vehicle on the left, and a value of 100 for the vehicle on the right.

Parameter 1, Belt Angle – The angle of rise of the beltline of the design from nose to tail, θ_1. This parameter affects the height of end points of Curves 6, and 7. Increasing this parameter makes the rear end of the vehicle taller, which gives the vehicle a more wedge-like shape. The effect of Parameter 1 is shown in Figure 3.2.

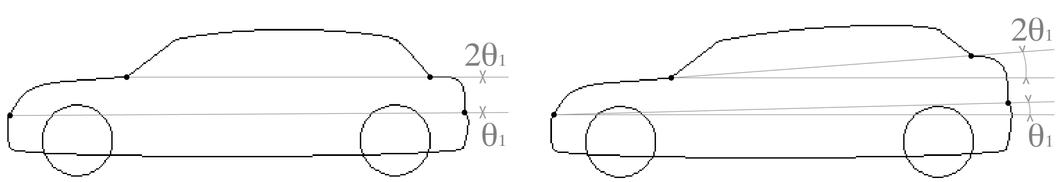

Figure 3.2 – The effect of changing Parameter 1, which affects the belt angle of the vehicle

Parameter 2, Nose Angle – The angle of rake of the nose of the car, θ_2. This parameter does not change the end point position of any curves, but changes two control points on Curve 2 to change the tangent angle of the leading edge of the hood. Increasing this parameter increases how much the nose leans back. In older versions of this model, as used in Sections 4.1 and 5.2, the nose angle of the model was fixed. The effect of Parameter 2 is shown in Figure 3.3.

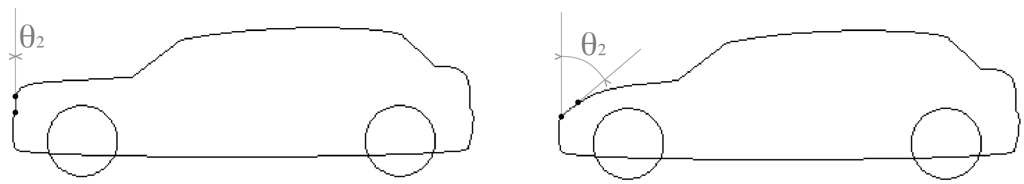

Figure 3.3 – The effect of changing Parameter 2, which affects the nose rake angle of the vehicle

Parameter 3, Ground Clearance – The distance from the floor pan to the ground, H_3. This parameter determines the height of the floor pan, which in turn affects the height of all curve end points by the same amount. The body shape is maintained, but is raised or lowered. Increasing this parameter increases the ground clearance. The effect of Parameter 3 is shown in Figure 3.4.

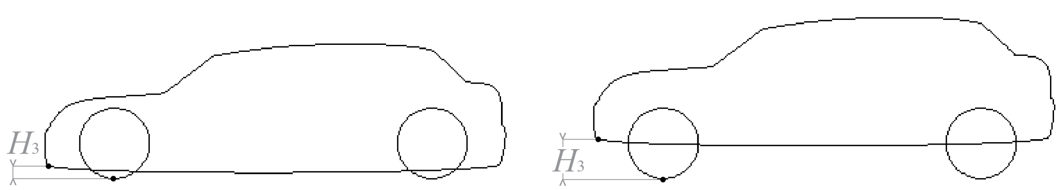

Figure 3.4 – The effect of changing Parameter 3, which affects the ground clearance of the vehicle

Parameter 4, Body Height – The height of the body of the car, H_4. The roof and windshields are not affected by this parameter. Increasing this parameter increases the height from the floor pan of the car and the top of the front bumper, and the height from the top of the front bumper to the cowl. The cowl is the point where the hood meets the windshield. Because the height of the rear bumper and the height of the trunk is determined from a combination between the belt angle and height of the front bumper and cowl, changing the body height also affects the height of the rear bumper and trunk height uniformly. This parameter affects the y-coordinates of Curves 1, 2, 6, and 7. In older versions of the model, as used in Sections 4.1 and 5.2, the body height is determined separately by two parameters, the trim height and the cowl height. The trim height determines the height of the bumpers, and the cowl height determines the height from the top of the front bumper to the cowl. It was determined from user feedback that some people had trouble determining the difference between the two heights since they both affected the height of the body of the car, so in later versions of the model, the height of the body is made up by changing these two previous parameters with a fixed ratio. With the reduction of parameter count due to this change, the rake angle of the nose could be added as described for Parameter 2. The effect of Parameter 4 is shown in Figure 3.5.

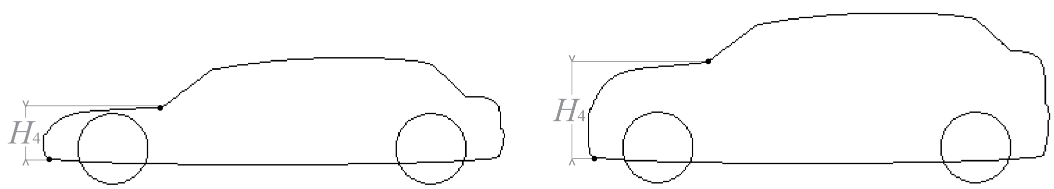

Figure 3.5 – The effect of changing Parameter 4, which affects the body height of the vehicle

Parameter 5, Roof Height – The height of the top of the windshield as measured from the cowl, H_5. This parameter affects Curves 3, 4, and 5. Increasing this parameter increases the height of the roof, both at the height of the windshield, and the curvature of the roof. Decreasing this parameter results in what is characterized as a chopped roof. The angle of the front and rear windscreens, as defined by Parameters 8 and 9 are unchanged, resulting in longer windows. The effect of Parameter 5 is shown in Figure 3.6.

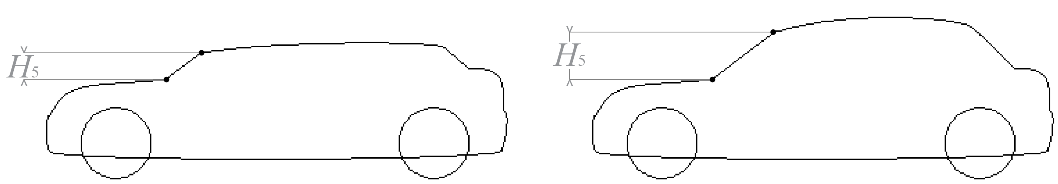

Figure 3.6 – The effect of changing Parameter 5, which affects the roof height of the vehicle

Parameter 6, Hood Length – The length of the hood, L_6, which is determined by the length from the back of the front bumper to the cowl, or base of the front windshield. This parameter affects Curves 2, 3, and 4. The overall profile of the hood is maintained, but is stretched back horizontally. Increasing this parameter increases the length of the hood. The effect of Parameter 6 is shown in Figure 3.7.

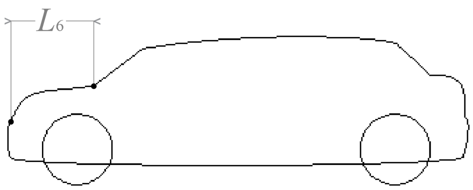

Figure 3.7 – The effect of changing Parameter 6, which affects the hood length of the vehicle

Parameter 7, Trunk Length – The length of the rear trunk, L_6, as defined as the distance from the back of the rear windshield to the front of the rear bumper. This parameter affects Curves 4, 5, and 6. Setting this parameter to a low value results in a station wagon, SUV, or hatchback style profile to the rear end to the vehicle, while setting this value higher results in a sedan style vehicle body. The effect of Parameter 7 is shown in Figure 3.8.

Figure 3.8 – The effect of changing Parameter 7, which affects the trunk length of the vehicle

Parameter 8, Front Windshield Rake Angle – The angle the front windshield leans backwards, θ_8. This parameter affects Curves 3 and 4. Increasing this parameter increases the rake angle of the front windshield. The height of the front windshield is defined by Parameter 5, which defines the roof height. Since the overall height of the front windshield is defined elsewhere, increasing this parameter also increase the length of the glass. The effect of Parameter 8 is shown in Figure 3.9.

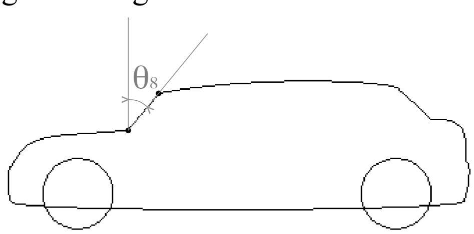

Figure 3.9 – The effect of changing Parameter 8, which affects the front windshield rake angle of the vehicle

Parameter 9, Rear Windshield Rake Angle – The angle the rear windshield leans forward, θ_9. This parameter affects Curves 4 and 5. Increasing this parameter increases the rake angle of the rear windshield. The height of the rear windshield is defined by Parameter 5, which defines the roof height. Since the overall height of the rear windshield is defined elsewhere, increasing this parameter also increase the length of the glass. The effect of Parameter 9 is shown in Figure 3.10.

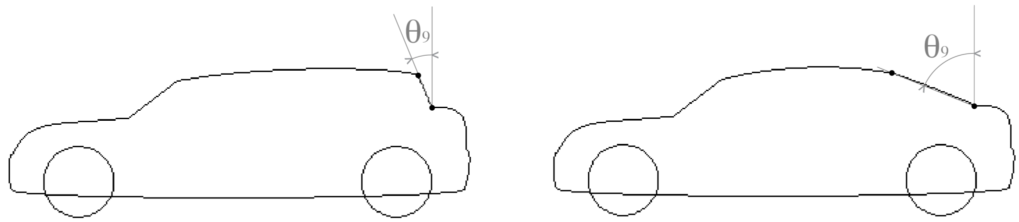

Figure 3.10 – The effect of changing Parameter 9, which affects the rear windshield rake angle of the vehicle

Parameter 10, Wheel Size – The diameter of the wheels on the vehicle, d_{10}. This parameter only affects the diameter of the circles, and does not affect any of the curves. The front and rear wheels are changed in size together, and the ground clearance of the vehicle is maintained. The effect of Parameter 10 is shown in Figure 3.11.

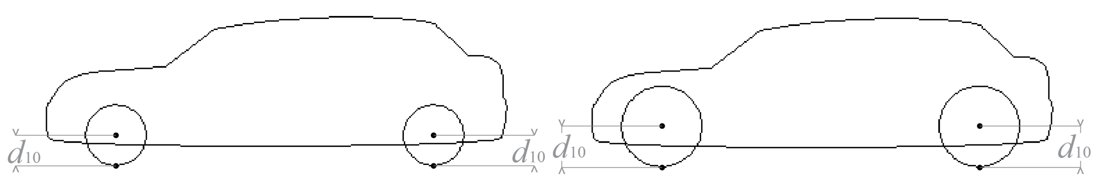

Figure 3.11 – The effect of changing Parameter 10, which affects the size of the wheels on the vehicle

Parameter 11, Front Wheel Position – The length distance of the front wheel from the front edge of the floor pan, L_{11}. This parameter only affects the positioning of the front wheel circle, and does not affect any of the curves. Increasing this parameter brings the front wheel rearward. The effect of Parameter 11 is shown in Figure 3.12.

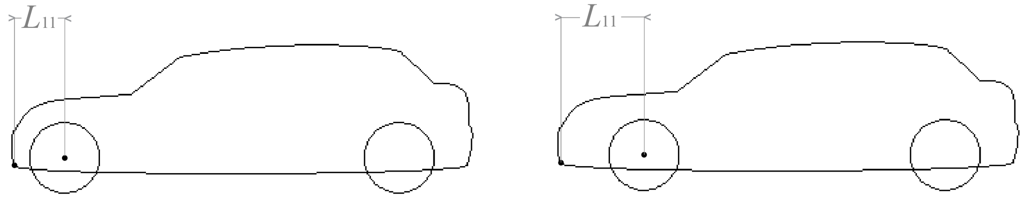

Figure 3.12 – The effect of changing Parameter 11, which affects the front wheel position of the vehicle

Parameter 12, Rear Wheel Position – The length distance of the rear wheel from the front edge of the floor pan, L_{12}. This parameter only affects the positioning of the rear wheel circle, and does not affect any of the curves. Increasing this parameter brings the rear wheel rearward. The effect of Parameter 12 is shown in Figure 3.13.

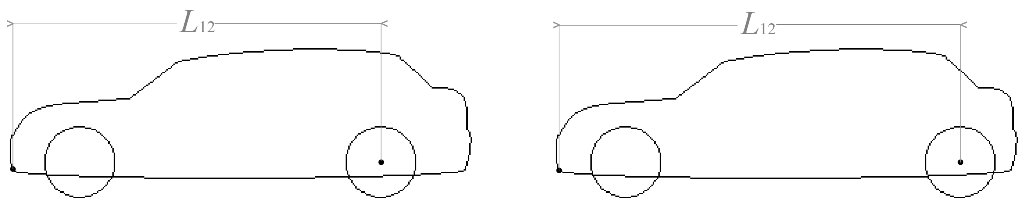

Figure 3.13 – The effect of changing Parameter 12, which affects the rear wheel position of the vehicle

3.1.2 DEFINING CURVES FROM PARAMETER VALUES

Each of the points used in this parametric car model are defined in a specific order, building the vehicle from the front to the rear, and from the bottom up. The reason for this order is that curves higher up on the vehicle are dependent on the position of curves below them for continuity. The roof is drawn last because it is dependent on all of the curves on the rest of the vehicle. The curve numbers presented in this section refer to Figure 3.1, and the parameters refer to the parameters described in Section 3.1.1. The plot dimensions for the vehicle design is 900 units wide and 500 units tall. In discussing the relationship between the curves and the parameter values, all of the parameter values are linearly scaled from a 0-100 range to the units used in the model. The scaling also defines the upper and lower bounds modeled for each of the parameters based on ranges observed in production vehicles and to ensure model stability. The MATLAB code for this parametric vehicle design model can be found in Appendix A.

The first curve that is defined is Curve 8, the floor pan of the car, and is drawn from front to back. Because the length of the vehicle is constrained, the only parameter needed to draw the floor pan of the vehicle is Parameter 3, which defines the ground clearance. Increasing the value of Parameter 3 increases the vertical height of the floor pan from the origin. The front endpoint of the floor pan is defined based on a fixed x-coordinate from the origin, and the y-dimension of front endpoint is defined based on a linear scaling of Parameter 3, which defines the ground clearance. The rear endpoint of the floor pan is defined by adding the fixed chassis length of the vehicle to the x-coordinate of the front end point of the floor pan. The two remaining control points are defined in fixed relation to the front and rear points to result in a slight curve and transitions for the front and rear bumpers.

The second curve that is defined is Curve 1, the front bumper of the car, and is defined from bottom to top. The lower endpoint of the front bumper is assigned to the same point as the front endpoint of Curve 8, the floor pan, for continuity. The upper endpoint of the front bumper is found by adding the trim height to the y-coordinate of the lower endpoint. The trim height is defined using a linear scaling of Parameter 4, which defines the body height. The two control points for Curve 1 are defined in reference to the endpoints of Curve 1 to resemble a typical vehicle front bumper for the full range of Parameter 4, and to maintain continuity with the control point for the front edge of Curve 8.

The third curve that is defined is Curve 2, the hood of the car, which is defined from front to rear. The front endpoint of the hood is assigned to the same point as the top of the Curve 1, the front bumper, for continuity. The x-coordinate of the rear endpoint of the hood is defined by adding the hood length, which is a linear scaling of Parameter 6, to the x-coordinate of the front endpoint. The y-coordinate of the rear endpoint of the hood is defined by adding the cowl height, which is a linear scaling of Parameter 4, to the y-coordinate of the front endpoint of the hood. The control points of Curve 2 are defined by the endpoints of Curve 2, and Parameter 2, which defines the nose rake angle, such that a parameter value of 100 yields a 45 degree nose rake angle and a parameter value of 0 yields a vertical nose rake angle.

The fourth curve that is defined is Curve 3, the front windshield, which is defined from bottom to top. The bottom endpoint of the front windshield is assigned to the same point as the rear of Curve 2, which defines the hood. The y-coordinate of the front windscreen is found by adding the roof height, which is a linear scaling of Parameter 5, to the y-coordinate of the bottom endpoint of the front windscreen. The front windscreen rake angle is a linear scaling of Parameter 8. The x-coordinate of the front windscreen is calculated using trigonometry as the roof height multiplied with the tangent of the front windscreen rake angle. The control points of Curve 3 are set equal to the endpoints, so that the windshield will not have any curvature from the side profile, which is true for most production cars.

The roof cannot be defined at this time because the position where it will meet with the rear windscreen is still unknown. So, the fifth curve that is defined is Curve 7, the rear bumper, which is defined from top to bottom. The bottom endpoint, control points, and x-coordinate of the upper endpoint of Curve 7 is defined similarly to how they were defined on Curve 1 using their continuity with the rear of Curve 8, the floor pan. The y-coordinate of the upper endpoint of the rear bumper is defined as a relationship between the trim height, which is a linear scaling of Parameter 4, the belt angle, which is a linear scaling of Parameter 1, and the chassis length, which is fixed. This y-coordinate is found by adding the trim height and the chassis length multiplied with the sine of the belt angle to the y-coordinate of the bottom of Curve 7, the rear bumper. This relationship results in the rear trim height being higher than the trim height in front based on the rising belt angle. In the case where Parameter 1 is 0, the belt angle is zero, in which case the rear trim height is the same as the front trim height.

The process used to determine Curve 2, the hood, and Curve 3, the front windscreen, is repeated for Curve 6, the trunk, and Curve 5, the rear windscreen, with two exceptions. The first exception is that in both cases is that the y-coordinates of their upper endpoints have an additional component due to the belt angle, which is calculated using the same method as described for Curve 7. The second exception is that the top of the rear windscreen is curved for a smoother transition with the roof, as is typical in most production cars.

The last curve to be drawn is Curve 4, the roof, which is drawn from front to rear. The endpoints of the roof are the same as the upper endpoints for the front and rear windscreens. The control point defining the curvature of the front of the roof is defined based on the length of the roof, which is found using the coordinates of the rear of the roof to allow for an appropriate curvature to integrate with the front and rear windscreens. The control point defining the curvature of the rear of the roof is based on the rear windscreen rake to maintain continuity of curvature where the rear of the roof meets the rear windscreen.

After all eight curves are drawn; the two circles for the wheels are drawn based on the wheel size, which is a linear scaling of Parameter 10, and the front and rear position of the wheels, which are found with a linear scaling of Parameters 11 and 12 respectively.

For more information about this parametric vehicle design model, please see the MATLAB code in Appendix A.

3.1.3 SAMPLE CAR DESIGNS

By setting values of the design parameters described in Section 3.1.2, the following sample of car designs can be generated as a vector of twelve numbers. The parameter values of these vehicles have been selected to generate vehicles that resemble a variety of real world vehicles, a 2009 Chevrolet Corvette as shown in Figure 3.14, a 2006 Toyota Prius as shown in Figure 3.15, and a 2010 Range Rover Sport as shown in Figure 3.16. The chromosome of each vehicle design is shown beneath the design.

These designs are also used to verify user consistency in later experiments, and are used in Section 3.2 to demonstrate correlations between the results from the aerodynamic model and the actual aerodynamic performance of the vehicles these profiles are based on. The parametric design model described in this section is specific to vehicle design, but any product category where designs can be concisely modeled using a similar parametric method can be substituted for use with the methods described in this dissertation.

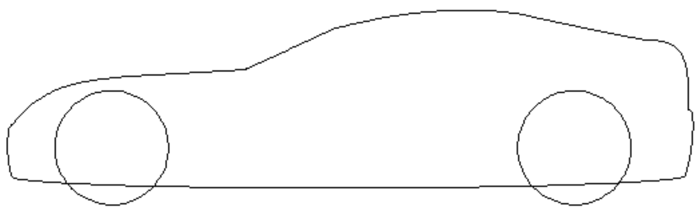

	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	
[	**90**	**60**	**20**	**50**	**0**	**100**	**40**	**95**	**100**	**80**	**50**	**30**	]

Figure 3.14 – Sample profile and chromosome designed to resemble a 2009 Chevrolet Corvette

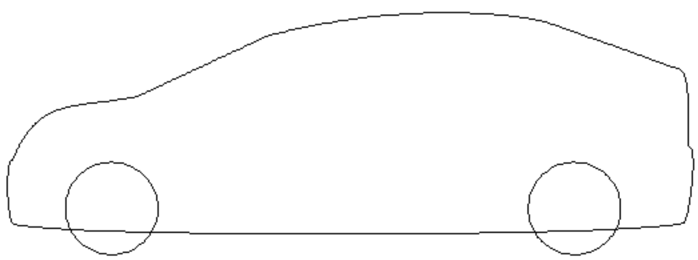

	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	
[	**90**	**70**	**30**	**80**	**70**	**0**	**10**	**95**	**100**	**10**	**50**	**30**	]

Figure 3.15 – Sample profile and chromosome designed to resemble a 2006 Toyota Prius

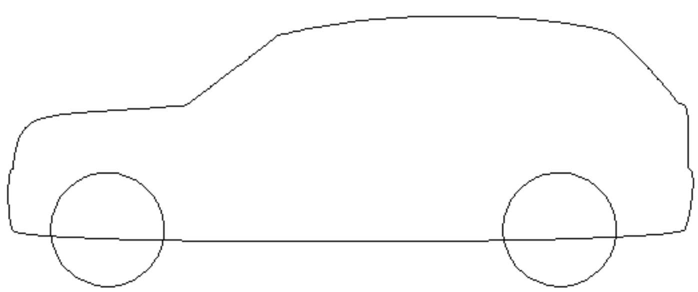

	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	
[	**5**	**15**	**90**	**78**	**100**	**50**	**5**	**50**	**50**	**80**	**40**	**0**	]

Figure 3.16 – Sample profile and chromosome designed to resemble a 2010 Range Rover Sport

3.2 AERODYNAMIC MODEL

Given the ever-increasing need for improved fuel efficiency, the issue of aerodynamic performance of vehicles has never been more important. Because of the rich interaction between the stylistic form of a vehicle and its effect on the vehicle's aerodynamic performance, this comparison is of utmost interest and is used throughout this dissertation as a functional performance assessment for the vehicles being designed. An experiment discussed in Section 4.1 examines vehicle design characteristics that communicate fuel efficiency to consumers. The model discussed in this section estimates the actual aerodynamic coefficient of drag based on the vehicle's chromosome, the parametric values of the vehicle's design.

In current production vehicles, the task of determining the coefficient of aerodynamic drag is typically performed in one of two ways, wind tunnel testing and Computational Fluid Dynamics (CFD) analysis. Wind tunnel testing, which is the most accurate method, uses a physical model or mockup of the vehicle and tests its actual drag coefficient in a wind tunnel. CFD analysis is branch of fluid mechanics that uses computational algorithms to model and approximate fluid flows around a body. CFD analysis can often be computationally very expensive (Becker et al., 2007). Due to the quantity of models to be assessed and the computational speed that is needed by the experiments, neither of these methods is suitable for use with this research. Instead, a feature based aerodynamic drag coefficient metamodel called CDaero by Guan (1995) was adapted for use with our parametric vehicle model. Note, CDaero was further refined and updated by Chan (Chan, 1997; Calkins & Chan, 1998), but for unknown reasons, this updated model yielded inconsistent results after being adapted for use with the parametric model. As a result, the model used in this dissertation is the model documented by Guan.

CDaero is based on work by Pershing and Masaki (1978), who developed eleven discrete parametric equations that modeled the primary contributions of aerodynamic drag as shown in the center column of Table 3.1. The coefficient of drag for a vehicle design is the summation of each of the eleven primary contributions. This model was updated by Carr and Stapleford (1981) for improved performance by identifying and individually treating different body feature types using the 13 primary contributions shown in the right column of Table 3.1. This model was updated by Guan (1995) for use with modern vehicle shapes, but retaining the same 13 primary contributions used by Carr and Stapleford, as pictured in the right column of Table 3.1. Guan's model uses 51 parameters to define vehicle shape and to calculate the 13 primary contributions to drag coefficient, and in test cases is able to predict vehicle drag coefficient to within +8.2% to -15.2% of actual wind tunnel results. This model was further refined by Chan (1998) to yield +/-6% accuracy on the same test cases.

Table 3.1 – 13 Primary Contributions to Aerodynamic Drag

Contribution to cD	Pershing and Masaki (1978)	Carr and Stapleford (1981) and Guan (1995)
cD1	Front End	Front End
cD2	Windshield	Hood
cD3	Front Hood	Windscreen
cD4	Rear Vert. Edge	Afterbody
cD5	Base Region	Skin Friction
cD6	Underbody	Underbody
cD7	Wheel and Well	Wheels
cD8	Rear Wheel Fairing	Wheel Wells
cD9	Protuberance	External Mirrors
cD10	Bullet Mirror	Drip-rails
cD11	Cooling	Window Recesses
cD12		Mudflaps
cD13		Cooling System

3.2.1 ADAPTING THE AERODYNAMIC MODEL

The vehicle model used in CDaero is more detailed and extensive than the parametric vehicle model developed for this research. Six of the 13 primary contributions, wheel wells, external mirrors, drip-rails, window recesses, mudflaps, and the cooling system were deemed extraneous to our needs since they modeled vehicle traits that did not exist in our model. Modeling only the seven remaining relevant contributions reduces the list of 51 necessary parameters to only 32 parameters. The vehicle model used in CDaero is three-dimensional, while our model is two-dimensional. Of the 32 parameters, 14 pertain only to width, or details that are irrelevant in our model, and thus were assumed to equal a constant average value determined from a 1994 Lexus LS400, which Guan (1995) used as his illustrating example due to its typicality and average dimensions. These simplifications result in a two-dimensional version of the CDaero model with only 18 parameters, which are in line with the level of detail needed in this research.

In order to convert our 12-parameter model into the 18 required parameters, analytical solutions for a number of heights, positions, and points were used. Remaining points that could not be ascertained analytically were found by computationally searching along various vehicle curves to find maximum curvature, points of inflection, and maximum and minimum height where relevant. Because of the linearity of the parametric model, the results from these searches at extreme parameter values allowed values to be approximated using interpolation. One such example was to determine the height of the roof, the maximum height of the Curve 4 in Figure 3.1 was found, by searching along the curve using De Casteljau's Algortihm, for the lowest and highest roof heights, and these values were linearly interpolated in the final model for the value at Parameter 5, which defines the roof height, to determine the roof height of the design. Because only seven of the thirteen contributions were modeled, and not all of CDaero's parameters were deemed relevant to our model, average fixed values for a 1994 Lexus LS400 were substituted for the remaining contributions and parametric values in CDaero. As noted above, this vehicle was chosen because it was determined by Guan (1995) to be an average vehicle to use as a baseline and illustrative example.

To better illustrate the adaptation process between the two models, the adaptation process for one of the seven contributions, C_{d4}, which models the effects of vortices on the afterbody of the car is described in greater detail here. C_{d4} is defined by Guan (1995) as shown in Figure 3.17. The definition for C_{d4} is provided below, and illustrated in Figure 3.18. The contribution of afterbody drag is dependent on eight variables, which were adapted as follows. All references to Curve numbers refer to Figure 3.1.

A_B is the projected 'base' area, which in our parametric model is the area from the bottom of the rear bumper to the height of the rear trunk lid. The coordinates of the bottom of the rear bumper can be located as the lower end point of Curve 7, and the coordinates of the top of the trunk lid can be approximated as the height of the top of Curve 6. As described earlier, due to the two-dimensionality of our model, the width of the vehicle is fixed. Multiplying the width of the car by the difference in heights results in the projected 'base' area.

H is the total height of the car, from the bottom or the chassis pan to the top of the roof. These values were found for extreme parameter values using De Casteljau's Algortihm and interpolated as discussed above.

A is the total projected frontal area of the car. This can be approximated by multiplying the total height of the car, as found above, with the average width of the car.

A_I is the projected inclined rear surface area, which is the projected area of the rear windscreen. This can be approximated by multiplying the difference in heights between the top of the rear windscreen and the bottom of the rear windscreen, which are known, by the average width of the car.

θ is the backlight slope angle, which is the angle of the rear windscreen, which can be found directly from Parameter 9, which defines the rear windscreen angle.

ϕ is the rear end step angle, which is the angle of a hypothetical line that connects the rear edge of the roof with the back edge of the trunk. This line can be made using the coordinates of the top of the rear windscreen and the peak of the trunk, which can be approximated as having the x-coordinate of the front of the rear bumper and the y-coordinate of the bottom of the rear windscreen. The angle of this line can be found as the arctangent of the ratio between the height and length of the line.

r_b is the mean radius of the trunk trailing edge, with a maximum value of $r_b/H = 0.12$. Since the parametric model does not vary this radius intentionally, and the radius used in the parametric model is large enough to always satisfy the maximum value ratio, $r_b/H = 0.12$ for all cases.

r_r is the mean radius of the trailing edge of the roof. This was found to vary mostly based on two variables, the roof height, and the angle of the rear windscreen. The radii for the four extremes were found graphically and were linearly interpolated for the model.

Substituting these variable values into the model yields the approximated afterbody drag contribution. Once summed with all the other contributions, the final approximated coefficient of drag is found.

$$C_{d_B} = 0.15\left\{\left(\frac{A_B}{A}\right) + \left(\left(\frac{C_{D_I}}{C_{D_B}}\right)_\theta + \left(\frac{C_{D_I}}{C_{D_B}}\right)_\phi\right)\left(\frac{A_I}{A}\right)\right\}$$

A_B Projected 'base' area

A Total projected frontal area of car

A_I Projected inclined rear surface area

For θ between 15 degrees and 41.5 degrees

$$\left(\frac{C_{D_I}}{C_{D_B}}\right)_\theta = 0.50 - -0.80\cos(4\theta)$$

For θ between 41.5 degrees and 90 degrees

$$\left(\frac{C_{D_I}}{C_{D_B}}\right)_\theta = 1.15 - 0.15\cos(4\theta)$$

For ϕ between 15 degrees and 41.5 degrees

$$\left(\frac{C_{D_I}}{C_{D_B}}\right)_\phi = 0.60 - \cos(4\phi) + 3.5\left[\frac{r_b}{H}\left(1 - \frac{w_o - w_b}{w_o}\right)^2 \sin(4\phi) - \frac{r_r}{H}\right]$$

For ϕ between 41.5 degrees and 90 degrees

$$\left(\frac{C_{D_I}}{C_{D_B}}\right)_\phi = 3.5\left(\frac{r_r}{H}\right)$$

θ: Backlight slope angle

ϕ: Rear end step angle

r_r: Mean radius of roof trailing edge

r_b: Mean radius of trunk trailing edge

H: Overall body height

$\frac{r_b}{H}$ and $\frac{r_r}{H}$ have a maximum value of 0.12

$\frac{w_r}{w_o}$ is the percent of trunk trailing edge with radius r_b

Figure 3.17 – Equations for calculating afterbody drag from Guan (1995)

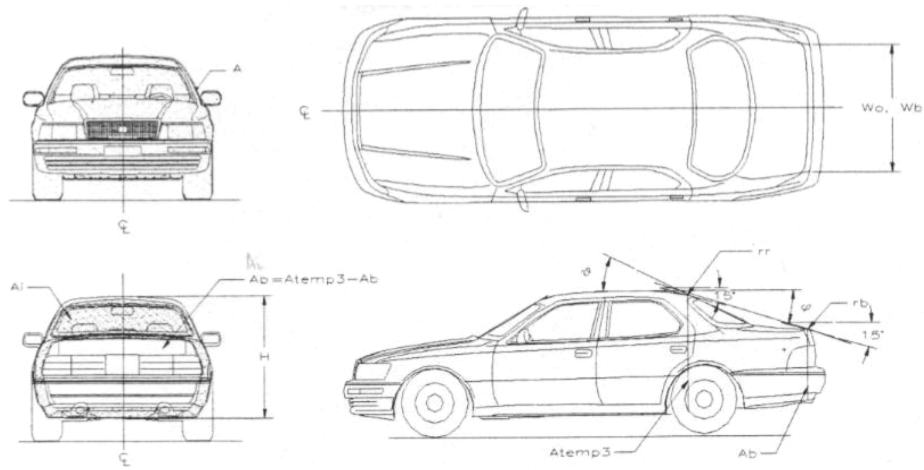

Figure 3.18 – Afterbody drag from Guan (1995)

The resulting model allows us to assign an aerodynamic drag coefficient to each vehicle design generated by our parametric model. While this estimated aerodynamic drag coefficient is not absolute, the trends of improved or worsened aerodynamic performance between different vehicle designs should be accurately captured. It should be noted that since this model outputs a coefficient of drag, a lower value corresponds to better aerodynamic performance.

3.2.1 MODEL VERIFICATION WITH ACTUAL CARS

The three vehicle designs shown in Section 3.1.2 and pictured in Figures 3.14, 3.15, and 3.16 are based on actual cars in production. The approximate coefficient of drag outputted by the aerodynamics model and the actual coefficient of drag by these production cars as found on Cars.com are shown in Table 3.2. The vehicles used in this comparison bear similarity to the actual cars, but do not represent the exact profile. Unsurprisingly due to the number of simplifications made to the CDaero model, the error rates are higher than those seen by Guan (1995) and Chan (1997), but do appear to show strong correlation with actual aerodynamic performance. For further verification of this Aerodynamics Model, please refer to Verification of Aerodynamics Model in Section 4.1.2.

Table 3.2 – Comparisons between actual cD and modeled cD

Vehicle	Actual cD	Modeled cD	Percent Error
2009 Chevrolet Corvette	0.28	0.244	12.8%
2006 Toyota Prius	0.26	0.279	6.8%
2010 Range Rover Sport	0.39	0.412	5.3%

3.3 VOLUMETRIC AND CENTER OF GRAVITY ASSESSMENT

Another functional assessment used in this dissertation to analyze vehicle designs is a measure of the volume of the vehicle and the center of gravity of the vehicle. It was suggested by Smith et al (2007), who were able to estimate the curb weight of vehicle designs with good accuracy by morphing different designs with known curb weights, that the correlations between shape and physical characteristics can be assumed to be related. While the usable interior volume of a vehicle is highly dependent on how the components in the vehicle are laid out, something that is not modeled with the current representation, it is expected that the volume of the body shape would be highly correlated with the total amount of usable volume in the vehicle, and thus would be used as one of the functional assessments. The height of the center of gravity of a vehicle can affect its handling ability and stability through difficult maneuvers, while the length coordinate of the center of gravity deals with whether a vehicle design may appear nose heavy or tail heavy in stylistic judgments. Also, consumers might base functional judgments on the center of gravity and total volume of the vehicle shape. Because the width of the vehicle is estimated to be fixed and constant, these calculations only find the area and center of gravity of the shape of the side profile of the vehicle, but are directly analogous with the volume and three-dimensional center of gravity in the two axes analyzed.

The following program was developed to calculate the volumetric and center of gravity properties of each car shape. The method of the program is analogous to numerical integration and is illustrated with the simplified example shown in Figure 3.19, which depicts shape AOF, a shape bounded by Bézier curve AF and line segments AO and FO. The area of shape AOF cannot be found analytically due to the complexities of curve AF. On the other hand, the area of shape AOF can be approximated by summing the areas of the five triangles within shape AOF, AOB, BOC, COD, DOE, and EOF, which can be found analytically. The vertices of the triangles along curve AF can be found by parsing the Bézier curve into six segments using De Casteljau's Algorithm. Similarly, the center of gravity of each individual triangle can be easily located analytically by averaging the coordinates of each triangle. By weighting each triangle by the area of the triangle, a weighted average can be easily constructed to determine the coordinates of the center of gravity of shape AOF. This method only works if all lines projected from point O to curve AF intersect with the inside of curve AF. By counting the number of intersections with curves for each ray, this problem can be compensated for, but a description of this method is outside of the scope of this dissertation.

In order to determine the area and center of gravity of each car shape, the program parses each of the eight Bezier curves used to draw the vehicle shape uniformly into 100 line segments, and uses these segments to form triangles around a central point known to be within the vehicle shape and below the top of the front bumper. This location is selected because none of the rays radiating from the central point will intersect with more than one curve of the vehicle design. The area of the resulting 800 triangles, which can be found analytically, can be summed to approximate the volume of the vehicle shape. Similarly, the 800 triangles can also be used to find the x and y coordinates of the center of gravity by summing an area-weighted average of the coordinates of each individual triangle. The resulting area is multiplied by the constant width of the car to find the volume.

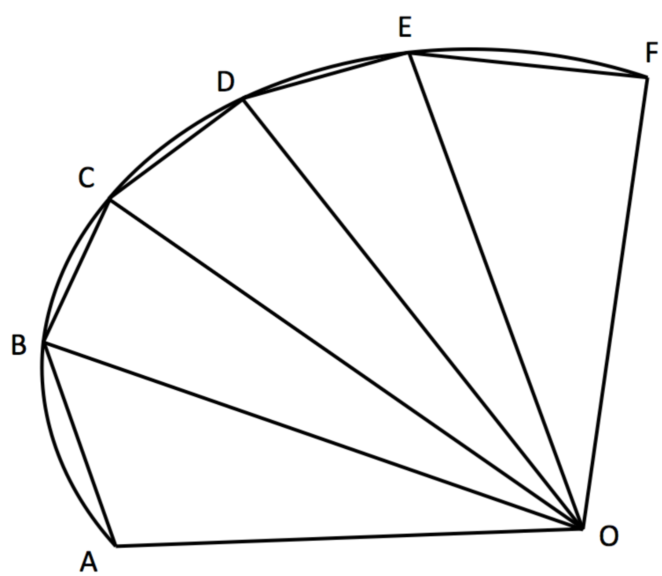

Figure 3.19 – Illustrative shape AOF

CHAPTER FOUR
UNDERSTANDING CONSUMERS

Two of the primary goals of the work presented in this dissertation are to create a computer model of consumer form-based shape judgments, and to examine ways to improve designers' ability to manage multiple conflicting goals in stylistic form and function. Before either of these goals can be approached, we need to gain a better understanding of the consumers we will be modeling, which is the focus of this chapter. This better understanding forms the foundational methods used to measure and record form-based consumer judgment for use in computer modeling in Chapter 5, and sheds light on the tradeoffs that exist between different stylistic form and functional goals that are studied in Chapter 6.

4.1 Form Function Fidelity

The relationships between different consumer judgments of stylistic form and function were studied and compared with actual measured functional performance for a better understanding of how human style and performance perceptions are created and how accurately they correlate with measured performance. Understanding these relationships allows for better management of tradeoffs, potential synergies in the design process.

This section sets out to better understand how people judge functional performance of products, and how well these judgments compare to indicators of actual performance. More specifically, participants were asked to rate how aerodynamic, sporty, fuel efficient, and rugged a computer generated car design appeared to them, and these ratings were analyzed against the actual aerodynamics of the vehicle as well as key indicators of sportiness and cornering stability such as the height, volume, and center of gravity of the vehicle and the vehicle's wheelbase. The inter-rater consistency between participants of human judgments was also studied. Using this human judgment data, the attributes in car design with the greatest effect on participant judgment of vehicle performance were identified, and were compared against their importance and effect in actual vehicle performance.

Reid et al. (2009) presented a methodology for quantifying the perceived environmental friendliness of vehicle silhouettes. In that paper, participants rated computer generated vehicle designs on how environmentally friendly the design appeared, or how likely the designs were inspired by nature. The data was correlated with the actual physical shape for each vehicle design. Reid et al. found that vehicles with shape discontinuities, leading to a boxier shape, were less likely to be perceived as inspired by nature, and that vehicles with more raked front and rear windscreens, and more gradual transitions into the roofline, in turn are more likely to be perceived as environmentally friendly.

The work presented in this chapter builds on the findings of Reid et al. by studying the effects of performance criteria on consumer perceptions. The performance criteria used are the actual aerodynamic performance, the height of the vehicle, the center of gravity of the vehicle, and the wheelbase of the vehicle. This work also extends Reid's work by studying a wider range of consumer perceptions, and by relating them to more elements of vehicle design, which are afforded by a more dynamic vehicle shape model.

4.1.1 METHODS

A computer-administered survey was set up in a booth at the 2009 Pittsburgh Vintage Grand Prix car show. Thirty-four attendees of the car show volunteered for this study. The data from eight participants, who were under the age of 18, or who did not successfully complete the survey, were removed from the data pool, resulting in twenty-six participants, eighteen male, and eight female.

4.1.1.1 Survey Data

In order to measure participants' judgments of each vehicle design, a computer administered semantic differential survey was used. The computer program showed each participant thirty computer-generated car designs in succession. The car designs were generated using an older version of the parametric generation model described in Section 3.1. The differences between the newer and older versions of the model are discussed in Section 4.1.1.2. The first five cars shown to each participant were the same across all participants to allow for inter-rater reliability tests between participants, and to ensure the same range of examples are initially shown to each participant to calibrate their responses for the remaining vehicles. The computer randomly generated the remaining twenty-five cars shown to each participant. Participants were asked to rate each car design on four criteria on a five-level Likert scale. The four criteria used were how well the words "Sporty", "Rugged", "Aerodynamic", and "Fuel Efficient" described each car design. A screenshot of the survey program interface is shown in Figure 4.1.

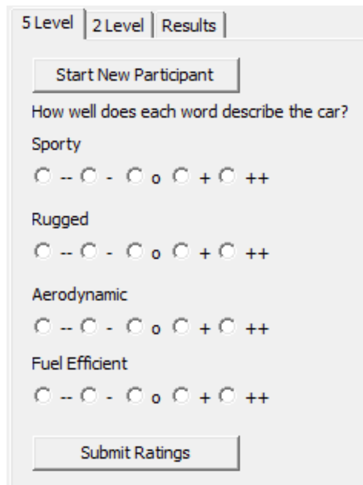

Figure 4.1 – Survey program interface

4.1.1.2 Parametric Design Generation

This study uses an older version of the parametric car design model presented in Chapter 3.1. The only difference is that this model uses a subtly different parameter list as shown in Table 4.1. In the model used in this study, nose rake angle is fixed, and cowl height and trim height are allowed to vary independently.

4.1.1.3 Functional Performance Models

The aerodynamic model described in Section 3.2 was used to assess the accuracy of each participant's ratings. The aerodynamic model outputs the coefficient of drag for each vehicle design, where a lower value indicates better aerodynamic performance. In the survey for this study, a higher value indicates better aerodynamic performance. For consistency with the results of human judgments of aerodynamics, the output of the model was inverted such that a better aerodynamics performance is indicated with a higher number.

Table 4.1 – Twelve design parameters in the vehicle chromosome

Gene	Design Parameter
1	Belt Angle – The angle of rise of the belt line from nose to tail
2	Ground Clearance – The distance from the floorpan to the ground
3	Trim Height – The top height of the bumper at the nose
4	Cowl Height – The height where the hood meets the windshield
5	Roof Height – The height of the top of the windshield
6	Hood Length – The length from the back of the bumper to the windshield
7	Trunk Length – The length from the back of the rear windshield to the front of the rear bumper
8	Windshield Rake Angle – The angle the windshield leans back
9	Rear Windshield Rake Angle – The angle the rear windshield leans forward
10	Wheel Size – The diameter of the wheels
11	Front Wheel Position – The length distance of the front wheel from the origin
12	Rear Wheel Position – The length distance of the rear wheel from the origin

The center of gravity and volumetric assessment model described in Section 3.3 was used to assess those functional characteristics of each vehicle design. In addition to the two functional performance models developed in Chapter 3, a measurement for wheelbase, which is the distance between the front and rear wheels, and a measure of the total height of the vehicle were also added to the analysis. These measures could be directly inferred from the chromosome of the vehicles.

4.1.1.4 Inter-rater Reliability

In order to assess the reliability and consistency of ratings between different participants on judging the same car, the first five vehicles presented to each participant were the same. To measure this inter-rater reliability, we chose to use intraclass correlations as documented by Shrout and Fleiss (1979). Specifically, we used a two-way random effects model with relative agreement because we're interested in extrapolating inter-rater agreement from the first five vehicle designs, thus we want to assume that the four judgment criteria (sportiness, aerodynamics, fuel efficiency, and ruggedness) are stable across judges. Inter-rater agreement (Cronbach's Alpha) was between 0.856 and 0.975, well above the typical threshold of 0.7 for all criteria, indicating that judgments were highly consistent across judges. See Table 4.2 for the actual alpha values and 95% confidence intervals for each of the four criteria. It is also interesting to note that participants were most consistent in judging the aerodynamics of car designs, followed by the sportiness, ruggedness, and were the least consistent in judging the apparent fuel efficiency of a car design.

Table 4.2 – Interrater Reliability

Criteria	Cronbach's Alpha	Confidence Interval (95%)
Aerodynamic	0.975	0.926 – 0.997
Sporty	0.962	0.889 – 0.995
Fuel Efficient	0.856	0.580 – 0.983
Rugged	0.908	0.732 – 0.989

4.1.2 RESULTS AND DISCUSSION

Pearson's product-moment correlation coefficients were calculated to measure the linear dependence between the human judgment survey data, functional performance data, and the each of the 12 vehicle design parameters in the vehicle chromosome. The resulting correlation matrices of this survey experiment can be found in Tables 4.3 and 4.4, and are discussed in the following sections. Table 4.3 is a correlation table comparing the twelve vehicle design parameters against measures of functional performance and human judgments of these vehicle designs. Similarly Table 4.4 is a correlation table comparing the measures of functional performance and human judgments against each other. In Tables 4.3 and 4.4, significance figures have also been listed for each comparison.

For the following results, significant correlations are deemed correlations with a significance value below 0.05, which are shown in the tables with a single asterisk and are shaded in lighter grey. Highly significant correlations are deemed correlations with a significance value below 0.01, which are shown in the tables with a double asterisk and are shaded in a darker grey.

4.1.2.1 Verification of the aerodynamics model

Studying the correlations between the vehicle design parameters and the parametric aerodynamics model can help to give insights about what vehicle design parameters are most important to improving vehicle aerodynamics, as well as to help verify the accuracy of the aerodynamics model. In the third row of Table 4.3, it can be seen that actual aerodynamics is highly correlated with the first five vehicle design parameters, which pertain to vehicle height. The aerodynamics model correlates a lower ground clearance, a lower trim height, a lower cowl height, and a lower roof height with improved aerodynamic performance. This agrees with common engineering knowledge, since a car with these traits would displace less air in motion and have a longer profile thus should have a lower coefficient of drag. Table 4.4 corroborates this by showing strong correlations between good aerodynamics and low center of gravity, low overall height, and low volume. The model also correlates a steeper belt angle and more rearward horizontal center of gravity with improved aerodynamic performance. These two traits both yield a more wedge-like shape to the car.

4.1.2.2 Human judgment of vehicle design parameters

Correlations between human judgments and vehicle design parameters are shown in Table 4.3. These correlations indicate that vehicles with low ground clearance, low trim height, low cowl height, longer hoods, and more steeply raked front and rear windscreens were judged more aerodynamic by participants. This follows common sense of what people commonly attribute to good aerodynamics. An example of a vehicle design with parameters in line with good aerodynamic judgments can be seen in Figure 3.14, which shows a parametric model designed to resemble a 2009 Chevrolet Corvette sports car.

Similarly, cars with even lower ground clearance, similarly low trim height, even longer hoods, similarly steeply raked front and rear windscreens and also large wheels were deemed to appear more sporty to participants. This also follows in line with traits common to sports cars, and is highly similar to the set of traits that communicate good aerodynamics. The same Corvette-like vehicle design, Figure 3.14, that is I line with good aerodynamic judgments is a good example of a vehicle design that would rank high in sportiness.

When judging fuel efficiency, participants preferred designs with low trim height and high roof heights, raked front and rear windscreens, longer trunk lengths, and rear wheels that are set farther forward. These traits tend towards designs of small modern family sedans, and the preference for longer trunk lengths tend away from silhouettes that resemble SUVs, vans, or station wagons. This finding agrees with Reid et al. (2009) who found that smoother roof transitions and less boxlike tail ends, which would be afforded by the higher roof height designs and longer trunk lids, would be perceived as more environmentally friendly.

Lastly, it was found that high ground clearance, high trim heights, high cowl heights, short hood lengths, and short trunk lengths led to higher scores in perceived ruggedness. These traits are indicative of typical SUVs, which tend to be tall, ride higher off the ground, have short hoods, and very short trunk lengths as shown by a vehicle modeled after a 2010 Range Rover Sport SUV in Figure 3.16.

By studying the relationships between human judgments of vehicle design parameters, it is suggested that consumers are in fact sensitive to varying vehicle traits and use them consistently to judge stylistic and functional form criteria. These relationships between human form judgment and vehicle design parameters will be modeled computationally in Chapter 5. It should be noted that these relationships discussed here are not an absolute definition of these stylistic and functional form criteria; rather they are an uncovering of some of the vehicle design traits that helped to elicit specific judgments from the participants in this study.

4.1.2.3 Human judgment of aerodynamics and other performance criteria

Table 4.4 describes correlations between human judgment and measures of performance criteria. Participants appeared to be good at judging the aerodynamic performance of vehicle designs. There was a strong positive correlation between the better aerodynamic performance and participant judgments of good aerodynamics. Similarly, vehicles that were rated as more sporty also tended to have better actual aerodynamic performance. Conversely, vehicles that were judged to be more rugged tended to exhibit poorer actual aerodynamic performance. Surprisingly, there was no measured correlation between judgments of fuel efficiency and actual aerodynamic performance.

In agreement with the correlations with aerodynamics, it can be seen that low vehicle height and low center of gravity are strongly correlated with high judgments in sportiness and aerodynamics, while high vehicle height and high center of gravity is strongly correlated with high judgments of ruggedness. This agrees with general engineering knowledge where lower vehicles have a lower center of gravity, which can translate into higher cornering stability, characteristic of a sports car, and cuts through less air when moving, which translates to better aerodynamics. Similarly, a taller vehicle design may afford improved ground clearance and interior volume, useful for rugged vehicle applications. It was also found that smaller car designs, or cars with lower volume, were found to correlate with high judgments in aerodynamics and sportiness, while larger volumes were found to correlate with high judgments in ruggedness.

On the other hand, neither the wheelbase of the vehicle designs, which is the distance between the front and rear wheels, nor the horizontal center of gravity of the vehicle designs appear to have any significant effect on human judgments of aerodynamics, sportiness, fuel efficiency, or ruggedness. This suggests that participants were insensitive to the wheelbase and horizontal center of gravity of the vehicles in forming their judgments. For this reason, these two functional measures were not revisited in later experiments.

4.1.2.4 Inter-relationships between human judgments of different dimensions

In order to better understand how people create form judgments of cars, the relationship between different human judgments must also be studied. The relationship between different human judgments can be found in the lower right corner of Table 4.4. One might expect strong positive correlations between the perceived aerodynamics of a car and the perceived fuel efficiency. This is in fact the case. More surprisingly, perhaps, is a strong positive correlation between the perceived sportiness and both perceived aerodynamics and perceived fuel efficiency. Most surprising, is a weaker but still significant positive correlation between the perceived ruggedness of a vehicle and the perceived fuel efficiency, but not the perceived sportiness or aerodynamics. This could perhaps indicate that the small SUV-like shape of some common hybrid vehicles may contribute to consumer perceptions that these shapes are more fuel efficient, despite not being perceived as aerodynamic.

Table 4.3 – Performance and human judgment correlations with vehicle design parameters

		Parameter 1 Belt Angle	Parameter 2 Ground Clearance	Parameter 3 Trim Height	Parameter 4 Cowl Height	Parameter 5 Roof Height	Parameter 6 Hood Length	Parameter 7 Trunk Length	Parameter 8 Windshield Angle	Parameter 9 Rear Window Angle	Parameter 10 Wheel Size	Parameter 11 Front Wheel Position	Parameter 12 Rear Wheel Position
Overall Height (higher=taller)	Correlation	-.005	**.520****	**.546****	**.538****	**.457****	.020	-.016	.049	-.045	.036	-.034	.018
	Significance	.892	**.000**	**.000**	**.000**	**.000**	.611	.676	.201	.244	.354	.377	.638
Wheelbase (higher=longer)	Correlation	.020	.003	.049	.030	-.006	.013	**-.079***	-.036	.032	-.024	**-.729****	**.704****
	Significance	.602	.946	.203	.434	.876	.746	**.041**	.351	.412	.538	**.000**	**.000**
Adjusted Actual Aerodynamics (higher=better)	Correlation	**.149****	**-.675****	**-.320****	**-.318****	**-.279****	.018	.015	.035	-.024	-.040	-.024	-.028
	Significance	**.000**	**.000**	**.000**	**.000**	**.000**	.641	.705	.372	.528	.301	.527	.471
Center of Gravity x (higher=rear)	Correlation	**.443****	**-.307****	**.114****	.050	**.177****	**.407****	**-.545****	**.113****	**-.375****	-.014	-.034	.053
	Significance	**.000**	**.000**	**.003**	.197	**.000**	**.000**	**.000**	**.003**	**.000**	.723	.375	.175
Center of Gravity y (higher=top)	Correlation	**.084***	**.686****	**.658****	**.354****	**.152****	**-.100****	**-.135****	.028	**-.080***	.044	-.031	.034
	Significance	**.030**	**.000**	**.000**	**.000**	**.000**	**.010**	**.000**	.466	**.038**	.254	.423	.385
Volume (higher=bigger)	Correlation	**.178****	**.672****	**.267****	**.586****	**.252****	**-.147****	**-.206****	-.027	**-.198****	.017	-.025	.034
	Significance	**.000**	**.000**	**.000**	**.000**	**.000**	**.000**	**.000**	.493	**.000**	.655	.511	.387
Judgment of Aerodynamics (higher=better)	Correlation	.005	**-.094***	**-.191****	**-.125****	-.008	**.134****	.001	**.142****	**.154****	.002	.016	-.043
	Significance	.896	**.015**	**.000**	**.001**	.841	**.001**	.982	**.000**	**.000**	.961	.680	.263
Judgment of Sportiness (higher=better)	Correlation	.042	**-.136****	**-.191****	-.062	-.043	**.231****	.034	**.141****	**.131****	**.136****	.047	-.036
	Significance	.280	**.000**	**.000**	.109	.264	**.000**	.386	**.000**	**.001**	**.000**	.229	.356
Judgment of Fuel Efficiency (higher=better)	Correlation	.039	-.003	**-.106****	-.002	**.152****	.046	**.079***	**.085***	**.106****	-.006	-.016	**-.082***
	Significance	.316	.947	**.006**	.960	**.000**	.238	**.041**	**.028**	**.006**	.883	.681	**.034**
Judgment of Ruggedness (higher=better)	Correlation	.012	**.203****	**.100****	**.142****	.024	**-.098***	**-.110****	-.026	-.022	.007	-.001	-.023
	Significance	.761	**.000**	**.010**	**.000**	.529	**.011**	**.005**	.496	.565	.853	.987	.556

Table 4.4 – Inter-relationships between performance and human judgment correlations

		Overall Height (higher=taller)	Wheelbase (higher=longer)	Adjusted Actual Aerodynamics (higher=better)	Center of Gravity x (higher=rear)	Center of Gravity y (higher=top)	Volume (higher=bigger)	Judgment of Aerodynamics (higher=better)	Judgment of Sportiness (higher=better)	Judgment of Fuel Efficiency (higher=better)	Judgment of Ruggedness (higher=better)
Overall Height (higher=taller)	Correlation	1	.037	-.776**	.013	.901**	.864**	-.202**	-.210**	.019	.229**
	Significance		.342	.000	.747	.000	.000	.000	.000	.616	.000
Wheelbase (higher=longer)	Correlation	.037	1	-.002	.060	.045	.041	-.041	-.058	-.045	-.015
	Significance	.342		.965	.119	.245	.290	.290	.137	.248	.695
Adjusted Actual Aerodynamics (higher=better)	Correlation	-.776**	-.002	1	.229**	-.758**	-.717**	.205**	.226**	.000	-.194**
	Significance	.000	.965		.000	.000	.000	.000	.000	.995	.000
Center of Gravity x (higher=rear)	Correlation	.013	.060	.229**	1	.001	.089*	.003	.065	-.047	.001
	Significance	.747	.119	.000		.978	.021	.942	.095	.226	.989
Center of Gravity y (higher=top)	Correlation	.901**	.045	-.758**	.001	1	.867**	-.238**	-.261**	-.056	.260**
	Significance	.000	.245	.000	.978		.000	.000	.000	.147	.000
Volume (higher=bigger)	Correlation	.864**	.041	-.717**	.089*	.867**	1	-.225**	-.241**	-.019	.282**
	Significance	.000	.290	.000	.021	.000		.000	.000	.621	.000
Judgment of Aerodynamics (higher=better)	Correlation	-.202**	-.041	.205**	.003	-.238**	-.225**	1	.568**	.359**	-.027
	Significance	.000	.290	.000	.942	.000	.000		.000	.000	.485
Judgment of Sportiness (higher=better)	Correlation	-.210**	-.058	.226**	.065	-.261**	-.241**	.568**	1	.163**	.059
	Significance	.000	.137	.000	.095	.000	.000	.000		.000	.129
Judgment of Fuel Efficiency (higher=better)	Correlation	.019	-.045	.000	-.047	-.056	-.019	.359**	.163**	1	.090*
	Significance	.616	.248	.995	.226	.147	.621	.000	.000		.020
Judgment of Ruggedness (higher=better)	Correlation	.229**	-.015	-.194**	.001	.260**	.282**	-.027	.059	.090*	1
	Significance	.000	.695	.000	.989	.000	.000	.485	.129	.020	

4.1.2.5 Principal Component Analysis

Principal Component Analysis is a mathematical technique that can be used to reduce the number of dimensions needed to describe a data set in which there are a large number of interrelated variables (Jolliffe, 2002). In this case it is a useful tool to determine the interrelations between different stylistic form judgments, and to analyze whether participants are actually bringing different judgments to the table. A look at the correlations between different stylistic form judgments in Table 4.4 indicates strong positive correlations between judgments of aerodynamics, sportiness, and fuel efficiency. There is also a weak positive correlation between judgments of ruggedness and fuel efficiency. Using Principal Component Analysis, two principal components with positive eigenvalues were identified and are shown in Table 4.6. These two principal components alone are able to explain 69.4% of the variation between the four judgment variables. The first principal component is comprised mostly of linear combinations of ratings of sportiness, aerodynamics, and fuel efficiency, and the second principal component is comprised almost entirely of ratings of ruggedness. What this means is that judgments of aerodynamics, sportiness, and fuel efficiency are highly interrelated and likely dependent on the same factors, and are highly independent to judgments of ruggedness. It is important to make the distinction that what is shown here does not indicate that ruggedness is the opposite of the other dimensions, but rather a wholly independent judgment that is not correlated with the trends of the others.

Table 4.5 – Two principal components that describe the design space

Component Matrix (2 components extracted)

	Component	
	1	2
Judgment of Aerodynamics	.874	-.174
Judgment of Sportiness	.792	-.116
Judgment of Fuel Efficiency	.596	.246
Judgment of Ruggedness	.101	.958

Extraction Method: Principal Component Analysis.

4.1.3 CONCLUSIONS

Whether it is due to years of product experience, advertising, or an innate ability to judge performance it appears that many consumers are able to accurately judge certain aspects of a car design's performance. Participants in this study were able to reliably and accurately gauge the aerodynamic performance of a variety of car designs. Participants were also able to reliably judge that car designs that are smaller, lower to the ground, and have a lower center of gravity are sportier and more aerodynamic, while larger and taller vehicles with a higher center of gravity were perceived to be more rugged. Perhaps this can all be explained by the literature (Parsons & Carlson, 2009; Ewen, 1990; Morgan, 1998; Harrison, 2001), which suggests that consumers often judge the functionality of a product using social norms and conventions of other products that the consumer is familiar with. As expected, it was found that the vehicle traits that best communicate specific perceptions of aerodynamics, sportiness, fuel efficiency, and ruggedness were generally in line with current vehicle classes that perform best on each of those criteria. The vehicle designs rated sportiest and most aerodynamic follow industry norms for sports car design, while designs rated most fuel efficient resembled common family sedans, and designs rated most rugged tended to resemble SUVs. On the other hand, consumer judgments appear not to be affected by the wheelbase or horizontal center of gravity of a vehicle design.

A potential warning against using social norms and conventions to educate consumers about the performance characteristics of a product is best illustrated by the lack of correlation between participant judgments of fuel efficiency and actual aerodynamics; two characteristics that engineering logic would expect to be strongly and positively correlated. Further examination reveals, vehicles that participants rated as more fuel-efficient tended to have profiles that resembled several popular hybrid vehicle models with taller proportions and higher roof heights. These shapes are selected for these hybrid vehicles due to a variety of utility, safety, comfort, and packaging reasons, but are not actually as aerodynamic as many lower-slung sports car shapes. In reality many of these sports car shapes, while very aerodynamic and thus having great potential to be fuel efficient, are often equipped with high power engines that negate any potential fuel savings. It is hypothesized that participants have been trained by the market to favor slightly taller rooflines as fuel-efficient over low-slung sports car like shapes, despite the latter vehicle design's aerodynamic advantage. This hypothesis is further reinforced by a positive correlation between judgments of fuel efficiency and judgments of ruggedness, which could perhaps indicate that the small SUV-like shape of some common hybrid vehicles may contribute to consumer perceptions that these shapes are more fuel efficient, despite not being perceived as aerodynamic. Further work needs to be done to examine this hypothesis and to better understand the lack of correlation between judgments of fuel efficiency and actual aerodynamic performance.

Principal Component Analysis helped to shed light on the relationships between different rated consumer judgments. It appears that out of the four categories of ratings that were surveyed, judgments of aerodynamics, sportiness, and fuel efficiency were highly interrelated, and ratings of ruggedness were largely independent. This lends insights towards experiments presented elsewhere in this dissertation by suggesting that it may not be necessary to survey all three of the interrelated categories. This finding also suggests that participants are not strictly rating preference with these vehicles, and possess at least two totally separate representations in these judgments.

In conclusion, this study indicates that consumers' form product judgments are based on their perceptions of performance, but that these perceptions are not always accurate. For example, simply designing a fuel-efficient car to perform well aerodynamically might not be sufficient to communicate the vehicle's fuel efficiency to consumers. As a result, product designers might be well heeded to investigate ways to better communicate a product's design and performance intentions and capabilities using form language that is appropriate and desirable for the product being designed. A method is developed in Chapter 5 that uses the survey techniques developed in this chapter to train an Artificial Neural Network and Genetic Algorithm based learning and generation system, which enables computers to model the relationships between human judgments and vehicle design parameters, and to computationally generate new designs that communicate the desired product design and performance intentions to consumers, while also incorporating the functional goals of the design. Furthermore, Chapter 6 examines how inspirational information can help designers to better manage the tradeoffs and synergies discussed in this chapter.

CHAPTER FIVE
MODELING CONSUMER JUDGMENTS

One goal of the work in this dissertation is to accurately model consumer aesthetic judgments of products with computational models, and to use these aesthetic models in concert with functional goals in a unified optimization design problem. The intent of the work documented in the previous chapter was to bring about a greater understanding of the relationship between aesthetic form judgments and expected performance of designs and to uncover a better understanding of how to model those consumer judgments.

In this chapter, three methods for modeling consumer judgments are explored. The first method represents the current state of the art for modeling consumer preference, logit modeling. In this case, logit modeling is used to gauge consumer preference towards functional features in laptop design. The second method discussed in this chapter uses Decision Tree Learning to model consumer aesthetic form judgments of computer generated car shapes. The third method extends the abilities shown with Decision Tree Learning by using Artificial Neural Networks to model consumer aesthetic form judgments. These Artificial Neural Networks were then coupled with numerical optimization tools to computationally generate vehicle designs that elicit targeted consumer responses and satisfy goals of both aesthetic form and function. This work makes no attempt to create an absolute definition of the terms being modeled; rather this work introduces a process that allows computers to uncover and model underlying factors behind human judgments of these terms, and to use these underlying factors to generate new designs that reflect these human judgments without the need of continued human interaction.

5.1 MODELING FUNCTION-BASED CONSUMER PREFERENCE USING LOGIT MODELS

The current state of the art in modeling consumer preference is the use of conjoint analysis. In order to better examine the capabilities and drawbacks of this technique, an experiment was conducted to gauge consumer preference for functional features in laptop computer design. More specifically, the method is similar to that of Michalek (2006), and is used here to gain a better understanding of the engineering tradeoffs in laptop design and to generate an optimal laptop design for college and graduate students. A logit model was developed based on data gathered from a conjoint survey to determine consumer preferences in the form of attribute splines. An integrated model was also created that ties the engineering model to the demand model to calculate the projected market shares of nine existing laptop models before and after the introduction of our optimal laptop. The optimal laptop design from this integrated model is found by minimizing the reduction in market share of the optimized laptop design. Due to the scope of this dissertation, the experiment overview in this section will only focus on the engineering model and the modeling of preference, and will not discuss the integrated model or the market share discussions. For more information about this study, please see Shiau, Tseng, Heutchy, & Michalek (2006).

5.1.1 ENGINEERING MODEL

An engineering model was developed that models key tradeoffs in laptop design. Each of the seven design variables, such as LCD screen size, various dimensions for the keyboard, body, and battery were bounded with realistic values from current laptop models and ergonomic constraints. A geometric representation of the dimensions used is shown in Figure 5.1. A number of engineering relationships were then inferred from a survey of current laptop models to estimate and constrain the cost, weight, battery life, and other laptop attributes based on the seven design variables. The four objective functions for this model are to maximize the LCD screen size, to maximize the battery life, to minimize the total weight, and to maximize size of the keyboard. The resulting optimization problem represents a continuous NLP model.

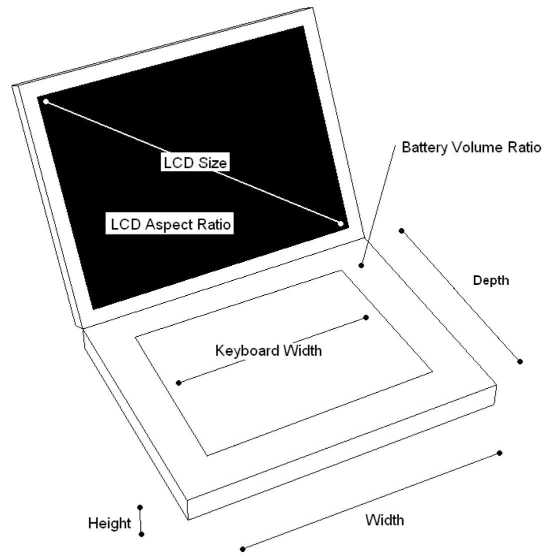

Figure 5.1 – Laptop engineering model

5.1.2 DEMAND MODEL

A demand model was then created using a discrete choice logit model, which used data from a conjoint survey. Survey-takers, who were told to assume that they are on the market for a new laptop, must choose one design that they would most likely buy out of the three choices, or "none" if they would choose none of the choices. The conjoint survey presented consumers with five attributes with which to judge the laptop designs, LCD size, thickness, battery life, weight, and price as seen in Table 5.1. Each value could take one of five levels as shown. It is important to note that the engineering model developed in Section 5.1.1 is not active for the survey. All designs, even infeasible designs, were presented to the survey-takers in an effort to model their preferences assuming anything was possible. The data used to create a logit model must be fully balanced for its results to be accurate. With five attributes and five levels, a full factorial survey would have required 3125 questions, which would certainly tax even the most patient survey-taker. SAS was employed to determine that only 25 questions were needed to keep the main effects clear and present a fully balanced partial factorial survey. One sample question is shown in Figure 5.2.

Table 5.1 – Attribute values in the conjoint analysis

Level	LCD size z_1 (inch)	Thickness z_2 (inch)	Battery life z_3 (hour)	Weight z_4 (lb)	Price z_5 (dollar)
1	10.4	0.75	1	2.5	750
2	12.1	1	2	4.5	1000
3	14.1	1.25	4	6	1250
4	15.4	1.5	6	8	1500
5	17.0	1.75	8	10	2000

1	A	B	C	None

LCD size	12.1 inches	17 inches	14.1 inches	-
Thickness	1.75 inches	1.5 inches	1 inch	-
Battery Life	4 hours	6 hours	2 hours	-
Weight	2.5 lbs	10 lbs	6 lbs	-
Price	1500 dollars	1000 dollars	2000 dollars	-
Width	11.3 inches	15.4 inches	13.0 inches	-
Depth	7.4 inches	10.0 inches	8.5 inches	-
	O	O	O	O

Figure 5.2 – Sample survey question

A standard logit model with discrete coefficients was employed as our demand model. Based on the standard logit model assumptions, the distribution of the unobserved error term ε_j was assumed to be an independently and identically distributed (iid) extreme value distribution (Train, 2003). The utility of a specific alternative can be expressed as a function of coefficients $\beta_{\zeta\omega}$:

$$v_j = \sum_{\zeta} \sum_{\omega} \beta_{\zeta\omega} \delta_{j\zeta\omega} ,$$

where j is the j-th alternative product, ω is the attribute level and $\delta_{j\zeta\omega}$ is chosen delta function respect to the j-th alternative with the ζ-th attribute at the ω-th level.

The logit probability P_j of j-th alternative is obtained by the following equation:

$$P_j = \frac{e^{v_j}}{\sum_k e^{v_k}},$$

where k is the number of all alternatives. The log-likelihood of logit model can be written as:

$$LL = \sum_j \sum_i \Phi_{ij} \ln(P_j).$$

The β coefficients can be calculated by solving the maximum log likelihood with optimization techniques.

5.2.3 ANALYSIS OF DISCRETE-COEFFICIENTS LOGIT MODEL

The discrete coefficients of each attribute are spline fitted, as shown in Figures 5.3, 5.4, and 5.5. Figure 5.3(a) shows the utility function of screen size. Unsurprisingly the highest utility is achieved at 15.4 inches, which is the most common laptop screen size currently on the market. As the screen size decreases, it appears that the utility of the design decreases roughly linearly. The utility also drops slightly for the larger 17 inch screen, likely because users deem the laptop too large to be portable at that point.

The general trend of utility of the laptop increases as the machine gets thinner and declines as the machine gets thicker as seen in Figure 5.3(b), the trend in between does not seem to make sense. Since the magnitudes of the part-worths are somewhat lower, showing weaker correlation, than other attributes, it can be assumed that with the exception of exceptionally thin or thick laptops, the thickness is less important to the surveyed consumers than other attributes. A larger survey sample would likely help to clear up this ambiguity.

Figure 5.4(a) shows the utility function in terms of battery life. The general trend is that customers prefer longer battery life, which agrees with intuition, but as battery life increases their sensitivity to increased battery life appears to reduce. This suggests that an increase in battery life matters more to the customer when the battery capacity is low, but once the capacity reaches a critical amount, which appears to be around six hours, the increase in utility tapers off. The effect of weight, shown in Figure 5.4(b), appears to be an approximately linear relationship. Users prefer light laptops to heavy ones; with an emphasized distaste for laptops over 8 pounds.

Figure 5.5 shows the effect of price on utility. There is a general preference for lower price, as is to be expected, but there is a flat plateau region in the middle of the plot indicating that users are less sensitive about the difference between $1,000 and $1,250 in price. The willingness to pay more than $1,500 seems to drop off rapidly with a drastically reduced utility at $2,000.

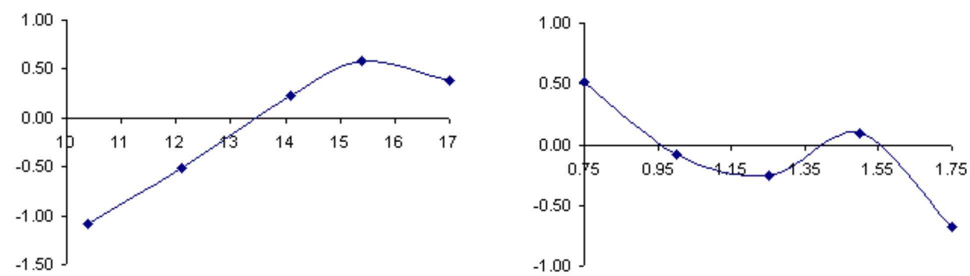

(a) (b)

Figure 5.3 – (a) Part-worth coefficient of LCD size (b) Part-worth coefficient of thickness

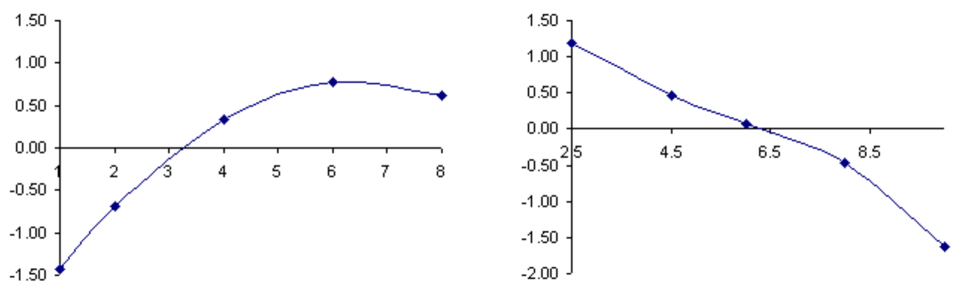

(a) (b)

Figure 5.4 – (a) Part-worth coefficient of battery life (b) Part-worth coefficient of weight

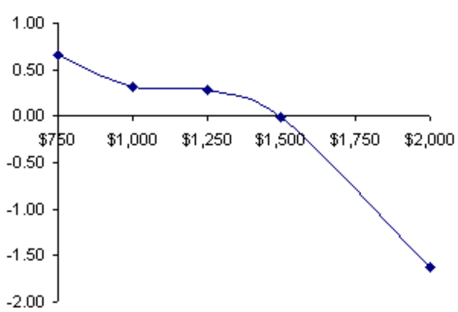

Figure 5.5 – Part-worth coefficient of the fifth attribute price

While this study focused on preference based on a combination of form and function considerations, the effects of stylistic form were not specifically targeted. Instead, the form being measured here was purely dimensional, and could be considered by many to be a functional requirement. As implemented in this study, this logistic regression model does not account for interaction effects between attributes. While this limitation may be acceptable in this case, since the attributes being examined are not highly dependent on each other, this limitation is unacceptable when modeling personal sensory opinions, such as stylistic and aesthetic form, where interactions between attributes play a more crucial role (Green & Srinivasan, 1990). It is possible to use conjoint analysis to model interaction effects using multiplicative interaction models (Friedrich, 1982; Aiken & West, 1991; Ai & Norton, 2003; Kam & Franzese, 2003), or by leveraging Non-geometric Plankett-Burman designs (Box & Meyer, 1993; Hamada & Wu, 1992; Blomqvist et al., 2007). Many common pitfalls and difficulties in modeling interaction effects using multiplicative interaction models have been identified and discussed by Brambor et al. (2005), who notes that in their survey of 156 articles that used multiplicative interaction models, only 16 (10%) implemented them correctly according to their checklist. Furthermore, the inclusion of interaction effects has been found to lead to models with lower predictive validity due to a larger number of model parameters (Green, 1984; Hagerty, 1986). Perhaps then it is not surprising that Wittink et al. (1994) found in a survey that only approximately 10% of commercial applications of conjoint analysis take interaction effects into account, while the majority focus only on the main effects.

Another limitation to conjoint analysis is the requirement of a balanced survey. This requires all survey takers be given a specific survey with a fixed number and layout of questions. This can prove difficult as the complexity of the model increases, and makes it nearly impossible to train the system on the fly. In order to address these problems, the next study uses decision tree learning to model stylistic preference.

5.2 LEARNING STYLISTIC FORM USING DECISION TREES

This section discusses a method that allows computers to learn human stylistic form judgments and preferences using Decision Tree Learning. Decision Tree Learning is a machine learning method that classifies training data by determining ways to split the data to gain the most information about the classifications. In this preliminary experiment, the decision tree receives its example input from surveys that record human judgments about a variety of randomly generated car designs. Using this decision tree, classifications can be made computationally to model the responses of people on new designs that were not part of the training set. To illustrate the abilities and drawbacks of the system, we have chosen to demonstrate this method on stylistic form ratings of vehicle design. The stylistic form judgments that will be learned by the computer are three criteria that are often used by humans to judge the appearance of a car: sportiness, ruggedness, and elegance. A fourth criterion of realism was also surveyed so the system could learn whether a car design appeared to be a realistic or feasible design. The methods discussed in this section can be used to computationally learn human stylistic judgments of a wide variety of product types, provided the product type can be represented using the methods discussed in Chapter 3.

5.2.1 SURVEY

In order to survey the aesthetic judgments of participants, a survey program similar to the one used in Section 4.1 was developed in C++ to conduct user surveys as seen in Figure 5.6. The developed survey program can conduct both two level and five level surveys, where the level number determines how many steps the user is given between the highest and lowest ratings in each category. In a two level survey, the user's choices are binary. Each car can only be ranked yes or no to questions of how sporty, rugged, elegant, or realistic the car is. In a five level survey, there are five choices ranging from '--' to '++'. The five level survey gives more detailed results due to higher resolution. The tradeoff is that the five level survey requires a significant amount more data to build a stable and descriptive decision tree. In cases where little data is available, the two level surveys are more robust.

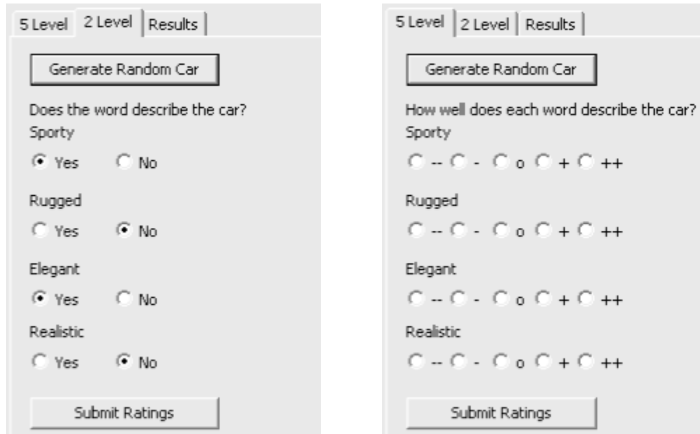

(a) **(b)**

Figure 5.6 **(a)** – Two level survey (filled out)
(b) – Five level survey (not filled out)

The survey program described above uses an earlier version of the parametric design model described in Chapter 3.1 to draw the cars presented to participants. There are two key differences between the two models. First, this model takes a three-level discrete input instead of a continuous input, and second, this model uses the same parameter list as used in Section 4.1, in that nose rake angle is fixed, and cowl height and trim height are allowed to vary independently. The twelve design parameters are represented as a vector of twelve integers, referred to as the vehicle design chromosome. Each digit can take the value of 0, 1, or 2, corresponding to three possible values each design parameter can take. These three values approximately represent high, medium, and low ranges of each parameter. The twelve design parameters and their definitions were shown in the previous chapter in Table 4.1.

5.2.2 DECISION TREE LEARNING

Decision trees have been used in data mining and machine learning for many years (Quinlan, 1990; Breiman, 1984). A decision tree is a predictive model that is constructed using categorized example data. The leaves of the completed decision tree represent the classifications, and the branches of the tree represent the differing parameter values that lead to the classifications. The method of the decision tree can be trained on any data where combinations of discrete input parameters yield discrete classifications. Further extensions have allowed decision trees to discretize on multiple intervals to allow the use of continuous valued data (Fayyad & Irani, 1992, 1993).

In the case of this example, four individual trees are constructed, one for each rated category, sportiness, ruggedness, elegance, and realism. A partial decision tree that classifies the sportiness of cars is shown in Figure 13 as an example. For each tree, the program begins by searching through the twelve design parameters to find which parameter will yield the highest information gain to split the data on. The information gain is based on the entropy and is calculated using the following formulas (Mitchell, 1997).

$$Gain(S,A) = Entropy(S) - \sum_{v \in Values(A)} \frac{|S_v|}{|S|} Entropy(S_v),$$

where

$$Entropy(S) = \sum_{i=1}^{c} - p_i \log_2 p_i,$$

where

p_i = the proportion of S belonging to class i, and

c = 2 for a two level survey and 5 for a five level survey.

Once the best split parameter is chosen, it is removed from a list of possible parameters to split the data on, and the total data is split for each of the three resulting branches, corresponding to the three values that each parameter can take. In Figure 5.7, it is assumed that the most information is gained by splitting the 62 initial examples of sporty cars and 47 initial examples of not sporty cars based on the cowl height of the cars. As can be seen after the split that vehicles with a low cowl height are far more likely to be sporty (46 cars) than not sporty (8 cars). The same trend in reverse is seen with high cowl heights.

Using only the subset of the data that falls under each branch, the system recursively begins again with each of the three branches. Each recursive loop continues until there are no more parameters to split on, the data is perfectly split, or there is no more data. As is seen in Figure 5.7, some cars can be perfectly classified with only two splits. In other words, if a car has a low cowl height and high wheel size, the car will be classified as sporty. In this example, all other branches must continue to be split using other parameters until one of the terminal conditions described earlier in this paragraph is met for each branch.

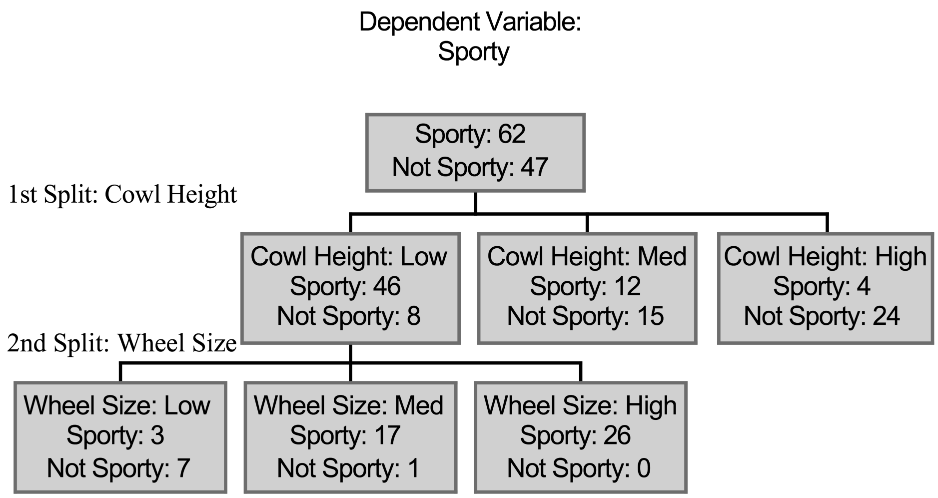

Figure 5.7 – Decision Tree Example

With larger amounts of training data, precautions have to be made to prevent overfitting the tree. A number of methods exist for neural networks and decision trees to prevent overfitting and to remove outliers from the data set (Mitchell, 1997; Quinlan, 1990; Breiman, 1984). Once the decision tree is built, any car chromosome can be fed to each of the four trees, for sportiness, ruggedness, elegance, and realism, and each tree will predict how the surveyed person would likely judge it.

5.2.3 RESULTS

Using a random 35 out of 53 two-level survey responses gathered from voluntary participants from a graduate level Mechanical Engineering class, four decision trees were constructed to model the sportiness, ruggedness, elegance, and realism of cars. These decision trees were then tested using the 18 remaining cars from the original survey responses to test for agreement between classifications from the decision tree and survey responses that were not used in the training process. The resulting percentage agreements are as follows. The decision tree was 78% accurate when modeling responses on whether a car appeared sporty, 78% accurate on whether a car appeared rugged, 61% accurate on whether a car appeared elegant, and 89% accurate on whether a car appeared realistic. The tree, on average, was 77% accurate at predicting the choices of the 18 participants. Another set of decision trees was then trained on all 53 two-level survey responses and was tested using a different set of 25 random cars that were generated using the survey program's test file function. All 25 test cars received classifications that were deemed appropriate by the experimenter.

Four illustrative cars were selected from the 25 test cars and are shown in Figures 5.8 – 5.11. According to the decision tree, Figure 5.8 was deemed sporty, not rugged, not elegant, and not realistic. Figure 5.9 was deemed not sporty, rugged, not elegant, and realistic. Figure 5.10 was deemed sporty, not rugged, elegant, and realistic. And Figure 5.11 was deemed not sporty, not rugged, not elegant, and not realistic.

It was hypothesized that adding more data points would dramatically improve the accuracy of the results. 434 two-level survey data points were collected from Mechanical Engineering graduate students to rate the sportiness of designs. The data was split such that 391 data points were used to train the decision tree and 43 points (10%) were set aside for validation. The resulting agreement was marginally improved from 77% to 79%. It was hypothesized that human inconsistency when making binary decisions about an abstract stylistic form judgment limits the maximum possible agreement, and that increasing the number of participants would not improve the results further.

The experiment was run again, but this time 114 five-level survey results were used to train the two-level decision tree. It was assumed the scale of the five levels corresponded to -2, -1, 0, 1, and 2 respectively. All negative results were lumped in one category, and all positive results were lumped in the other. All zero (middle) ratings were discarded, ratings of 1 and -1 were counted once, and ratings of 2 and -2 were repeated and used twice in the training of the decision tree, increasing their weighting in the training process. Using 100 data points for training and 14 for validation yielded an 86% agreement, a sizable improvement over the binary choice results.

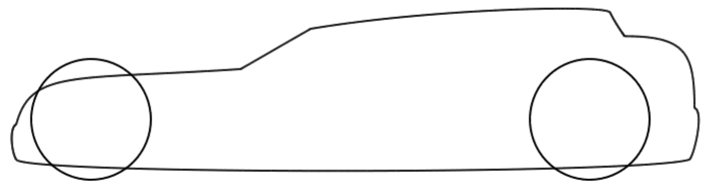

Figure 5.8 – Sample design deemed sporty, not rugged, not elegant, and not realistic by the decision tree

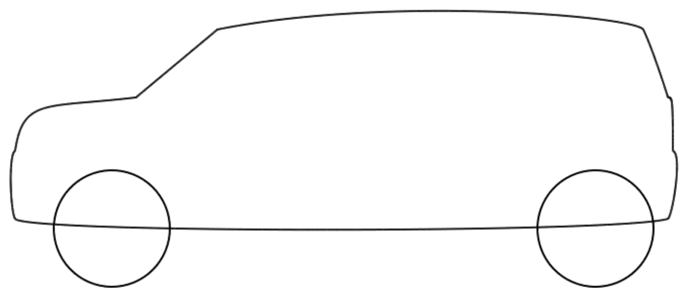

Figure 5.9 – Sample design deemed not sporty, rugged, not elegant, and realistic by the decision tree

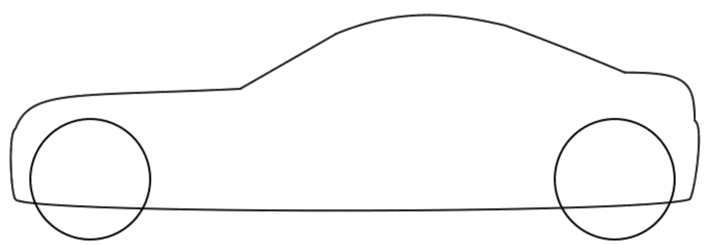

Figure 5.10 – Sample design deemed sporty, not rugged, elegant, and realistic by the decision tree

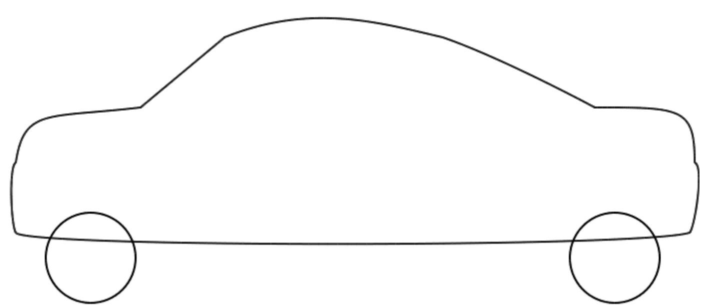

Figure 5.11 – Sample design deemed not sporty, not rugged, not elegant, and not realistic by the decision tree

5.2.4 DISCUSSION

Given that random assignment would produce approximately 50% agreement, these preliminary results are acceptable but not strong. It is important to keep in mind that this test is as much a test of the accuracy of the decision tree as it is the consistency of different participants, since there were numerous raters contributing to this data pool, and each participant is likely to have slightly different judgments. It appears that the participants were most consistent in their judgments of whether a car design appeared realistic, and the least consistent in their judgment of whether a car appeared to be elegant.

Unfortunately the accuracy of using more data points seemed to show little improvement in the accuracy of prediction, showing that human inconsistency is likely the limiting factor with subjective binary surveys such as this one. A majority of the two-level test cases where the decision tree did not match the actual survey results appeared to be in cases where the car design was borderline between classifications, for instance in the case of sportiness, the mismatched cases tended to be cars that were not strongly sporty nor unsporty when surveyed using a finer five-level scale of judgment. Using a five-level scale of judgment in the survey did seem to improve these numbers somewhat by removing more of the inconsistent data and double-counting the data that was higher in certainty. It is possible that with a larger data set and the described treatment of five-level survey data training a two-level survey could yield significantly better results.

In using decision tree learning instead of logistic regression analysis to model the data, the interaction effects between combinations of different parameters are more easily accounted for. Similarly, the robustness of decision trees to small and incomplete data is seen by the reasonable correlation of the results. Also, additional data can be added at will, and the system can work with as little or as much data as can be obtained. There is no need for the example data to be precisely balanced, as would be the case in logistic regression analysis (Green & Srinivasan, 1978).

Some drawbacks to using decision tree learning to model this data include the fact that the data is modeled as discrete instead of continuous. For instance, in the case of a five level survey, the decision tree does not understand that a rating of 4 is in between a rating of 3 and a rating of 5. All of the rating levels are deemed entirely independent. As a direct result, it would require a significantly larger set of data to run a full five level decision tree with the five-level survey data. A potential solution to this problem is to train a binary decision tree with five-level data, which incidentally resulted in the best results in this study. Another related drawback of decision trees due to their discrete nature is the difficulty of extracting trend directions. For instance, if someone were trying to make a car sportier, it would be difficult to directly extract whether the cowl height should be increased or decreased. The most serious ramification of this problem is that while the tree can classify a design based on training data, it is impossible to reverse the tree and extract parameter values from desired outcomes. This problem can be overcome by using numerical optimization or a genetic algorithm as discussed later in this chapter, but in order to best cater to these methods, the output of the form-learning algorithm should be continuous, whereas the output of a decision tree is natively discrete. The discrete nature of decision trees can be partially overcome by using regression trees, which accommodate continuous outputs (Breiman et al., 1984; Friedman, 1988; Quinlan, 2006). Nevertheless, in an effort to address these problems, the use of another machine learning method that is natively more robust to these problems, Artificial Neural Networks, is discussed in the next section.

5.3 LEARNING STYLISTIC FORM USING NEURAL NETWORKS AND GENERATING DESIGNS WITH GENETIC ALGORITHMS

Earlier in this chapter it was demonstrated that logit modeling can be used most easily to model consumer preference towards mutually independent features, and decision tree learning can be used to model discrete human stylistic form judgments. While work has been done to amend both of these drawbacks by modeling interaction effects in logit modeling (Friedrich, 1982; Aiken & West, 1991; Ai & Norton, 2003; Box & Meyer, 1993; Hamada & Wu, 1992; Blomkvist et al., 2007) and by adapting decision tree learning for continuous output by building regression trees (Breiman, 1984; Friedman, 1988; Quinlan, 2006), a method that is inherently able to combine the strengths of both logit models and decision trees in one aesthetic form model is desired. Specifically, an Artificial Neural Network (ANN) is able to natively model interaction effects between continuous features as well as provide a continuous output. One limitation to an ANN is that network inversion is impossible, so designs cannot be generated from the ANN directly. As seen in the literature (Baluja et al, 1994, Bull, 1999, Tsutsumi & Sasaki, 2008) it is common to use a Genetic Algorithm (GA) to perform network inversion, which allows the system to generate designs based on the ANN. In other words, the ANN, capturing people's aesthetic preference, becomes the fitness function for the GA, which can generate designs that optimize preference.

5.3.1 METHODS

A two-stage experiment was conducted where participants were asked to rate computer generated car profiles for sportiness, ruggedness, beauty, and fuel efficiency. In the first stage, participants generated survey data, which was used to train four ANNs per participant, one for each of the four rating categories. The resulting ANNs were then inverted using a Genetic Algorithm (GA) in order to generate new designs that the ANNs rate highest and lowest in sportiness, ruggedness, beauty, and fuel efficiency. In the second stage, the participants were surveyed with the cars generated using data from their ratings to verify the performance of the ANN and GA learning and generation system.

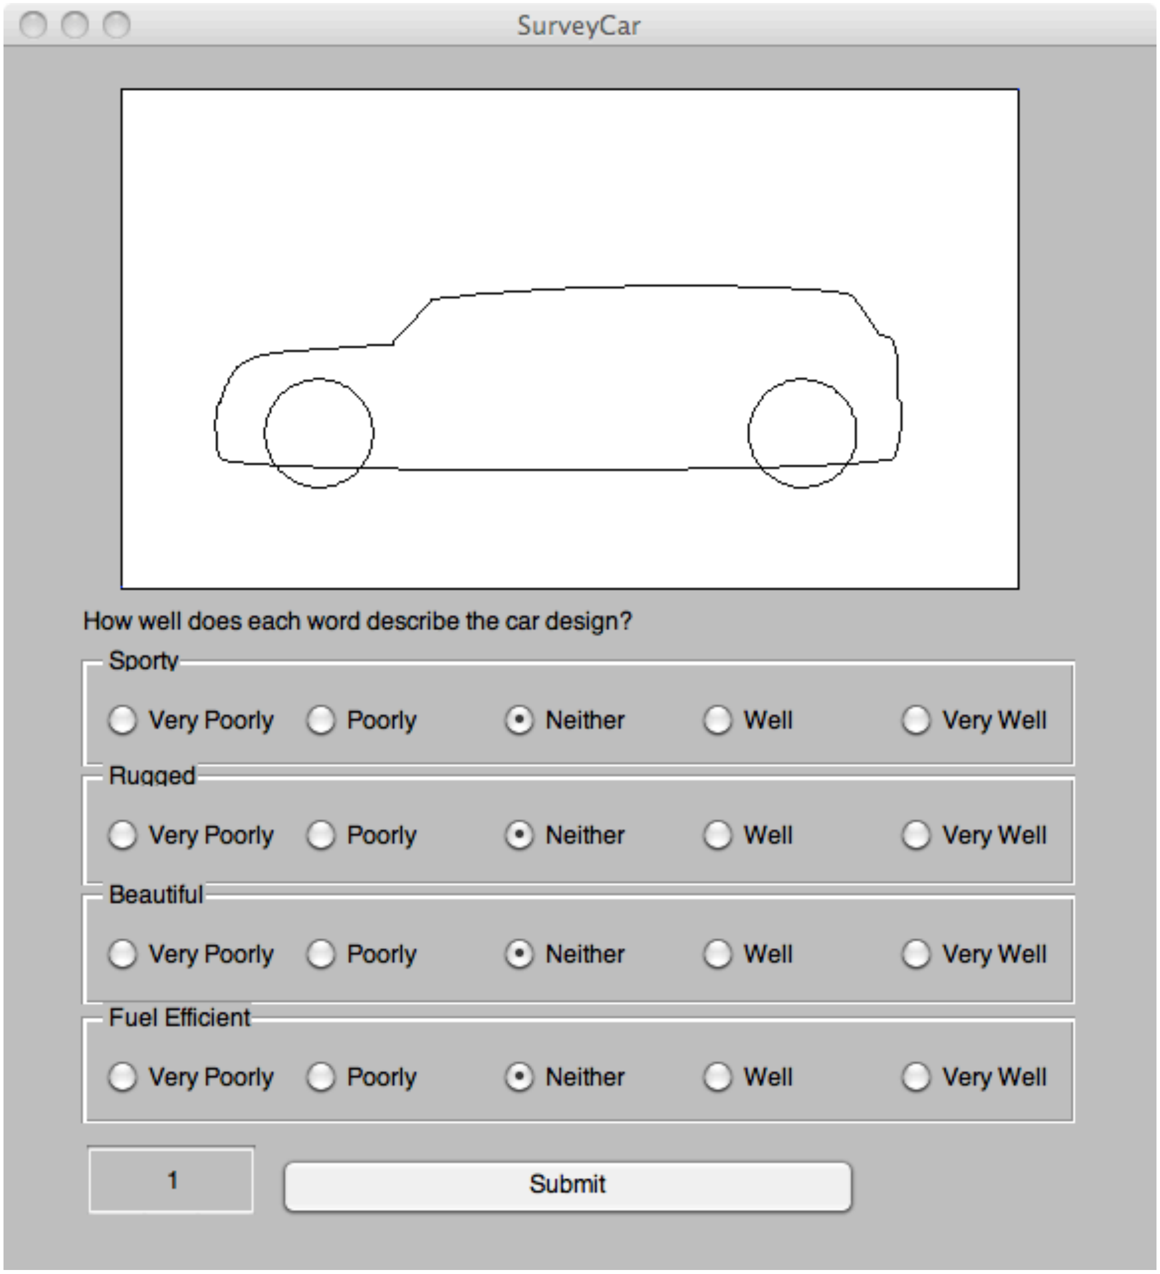

Figure 5.12 – Vehicle stylistic form judgment survey interface

5.3.1.1 Participants

Eighteen volunteers participated in this experiment ranging from 22 to 34 years of age, and a split of 11 males and 7 females.

5.3.1.2 Procedure

For both stages of the experiment, participants were presented with a computer interface written in MATLAB as shown in Figure 5.12. The interface displays a vehicle design generated using the parametric model detailed in Section 3.1 and asks participants to rank whether each of four words presented, Sporty, Rugged, Beautiful, and Fuel Efficient, described the vehicle design Very Poorly, Poorly, Neither, Well, or Very Well. Participants who were not already familiar with the range of vehicles that can be generated using this algorithm were first presented with a quick preview of 20 cars that were randomly generated using the algorithm. Participants were also encouraged to form a set of preference rules that govern the reasoning behind their judgments and to stick with them through both stages of the experiment to maximize consistency.

The first stage survey presented participants with 276 vehicle designs that were used to train the neural networks. This first stage survey took an average of approximately an hour for participants to finish, with the shortest taking 43 minutes, and the longest taking approximately two and a half hours to complete. Participants were encouraged to take breaks whenever they felt fatigued, with at least one break in the middle recommended. In between the two stages, the learning and generation system was trained on the data from the first stage of the experiment and verification designs were generated. Participants were asked to return for the second stage where they were presented with 50 vehicle designs to verify the performance of the learning and generation system. The break for participants between stages ranged from 3 hours to two weeks depending on scheduling restrictions with the median break being approximately 2 days. This second stage verification survey took an average of approximately 15 minutes, with the shortest taking 7 minutes and the longest taking 32 minutes.

5.3.3 DESIGN AND MATERIALS

For the first stage, each participant was presented with 276 computer generated vehicle designs to rate on a five-level scale in categories of sportiness, ruggedness, beauty, and fuel efficiency. The 276 designs presented consisted of 256 vehicles that were randomly generated in pre-assigned parameter ranges, and 20 cars inserted throughout for consistency verification. Details about the pre-assigned parameter range vehicle designs can be found in Section 5.3.3.1, and details about the consistency verification vehicle designs can be found in Section 5.3.3.2. The order of the 256 vehicle designs with pre-assigned parameter ranges was randomized, and then the 20 verification designs were inserted uniformly throughout the survey.

For the second stage, each participant was presented with a survey similar to the first stage, but with 50 computer generated vehicle designs. 24 of the 50 designs presented consisted of vehicles designed by the learning and generation system to elicit high and low ratings in each of the four categories. 12 of the 50 designs were the same cars inserted from the first stage for consistency verification. The remaining 14 of the 50 designs were randomly generated and were only presented to help disguise the 24 generated test designs. The order of the vehicles presented were randomized, and then any identical designs that appeared less than five questions apart were manually moved until no such violations were found.

5.3.3.1 Vehicle Designs with Pre-assigned Parameter Values

In order to ensure maximum diversity in the cars generated in the first stage of the experiment, parameter ranges were pre-assigned for 256 of the designs. The car design parametric model used here, as documented in Section 3.1, takes values between 0 and 100 for each of the twelve parameters. Parameters 1, 2, 4, 5, 6, 7, 8, and 9 in the model pertain to vehicle body shape, while Parameters 3, 10, 11, and 12 pertain to ground clearance, front and rear wheel position, and wheel size. Each parameter relevant to vehicle body shape was assigned a random value in either the upper half of the range (50-100), or the lower half of the range (0-50). To ensure that the population of cars presented in the survey has a maximum diversity of body shape, all upper and lower combinations of parameter values for parameters that pertain to vehicle body shape were presented in the survey, resulting in 256 car designs. The remaining four parameters were randomly generated in the full 0-100 range for these 256 cases.

5.3.3.2 Consistency Verification Vehicle Designs

In order to test for participant consistency, four predetermined car designs were presented five times each to each participant in the first stage, and three times each to each participant in the second stage, spread out uniformly throughout the surveys. The predetermined car designs were an average sedan, made up entirely of centered parameter values of 50, and the three vehicles designed in Section 3.1.2 to resemble a

2009 Chevrolet Corvette, a 2006 Toyota Prius, and a 2010 Range Rover Sport. The profiles for the sample cars can be seen in Figures 3.14, 3.15, and 3.16. This consistency verification data can help to determine how consistent a participant's ratings are with his or herself when the same car is presented multiple times, how consistent participants are between stages, and how much participants varied between each other. This verification data is used later to filter the results of the study.

5.3.3.3 Model Verification Vehicle Designs

To test the performance of the learning and generation system, designs were generated to target each participant's highest and lowest ratings in the four categories of sportiness, ruggedness, beauty, and fuel efficiency. This results in a total of eight vehicle designs. Tailor-made surveys were created for each participant using the high and low rating vehicles created using his or her own ANN. Each vehicle design was presented to the participant three times during the course of the second stage verification survey, resulting in 24 of the 50 designs presented. The order was randomized, and then filtered to ensure that no two identical cars could be presented within five cars of each other.

5.3.4 ARTIFICIAL NEURAL NETWORK

An Artificial Neural Network (ANN) is a computational model that contains a network of interconnected artificial neurons or nodes arranged in layers (Mitchell, 1997). As discussed in the background section, ANNs are able to make both discrete and continuous classifications and have strong abilities in pattern recognition, even with interaction effects between attributes (Fukushima, 1980, 1984). ANNs have been used successfully to perform character recognition (Fukushima, 2003), to recognize pleasing sounding phrases in jazz music (Griffith & Todd, 1999), and more recently to judge the aesthetic sensibility in roof structures (Tsutsumi & Sasaki, 2008).

An ANN consists of a network of artificial neurons or nodes arranged in layers as shown in Figure 5.13. Once created, the ANN is trained iteratively by processing data records one at a time and comparing the network's output with the training samples. The error at the output at each iteration is fed back into the network and is used to change the relative weights between the hidden nodes to improve network performance. This process iterates until no changes can be found to improve network performance. The data provided to the ANN is split up to perform a number of tasks. First, the bulk of the network is usually used to train the network. This is the data that is presented to the network during training to tune the weights of the network. Second, some of the data can be set aside for network validation. This data is not used directly to train the network, but is used to measure network generalization, and is used to stop network training if the results are starting to memorize specific instances in the training data rather than generalizing from their trends. Third, some data can be set aside to provide a testing set. This data has no effect on the training process and can be used as an independent measure of network performance. The number of nodes in the input and output layers are governed by the input data and desired targets, in the case of this study, each ANN contains 12 input nodes in the input layer and one output node in the output layer. There are no hard and fast rules for determining the number of hidden layers and hidden nodes in the hidden layer(s). Adding hidden nodes and layers generally increases the network's ability to model complex spaces and increases resolution, but if too many are added the ANN will simply memorize the training data and provide poor generalizability. A brute force method is employed in this work in an effort to circumvent this problem.

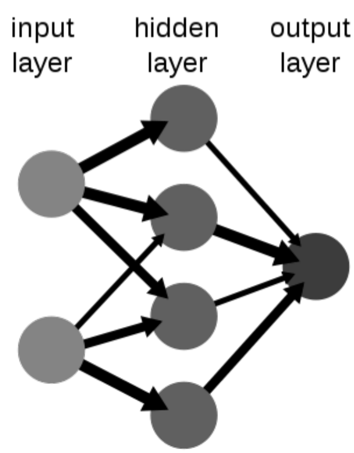

Figure 5.13 – A simple neural network

This research seeks only to use existing ANN technology to explore possibilities to unify stylistic form and function into a single framework. This research does not seek to advance the art of ANNs, so a readily available off-the-shelf implementation was used. The data in this experiment was used to train an ANN generated using the Neural Network Fitting Tool in MATLAB's Neural Network Toolbox. This tool creates a feed-forward network with one hidden layer with a variable number of hidden neurons. The network is trained using the Levenberg-Marquardt backpropagation algorithm, which is a numerical minimization algorithm for nonlinear functions and operates by interpolating between the Gauss-Newton and Steepest Gradient techniques. More information about this algorithm can be found in (Levenberg, 1944, Marquardt, 1963).

5.3.5 GENETIC ALGORITHMS

A Genetic Algorithm (GA) was selected as the method to invert the ANNs in this framework due to successes in the literature in coupling ANNs with GAs (Baluja et al., 1994; Bull, 1999; Tsutsumi & Sasaki, 2008). A GA is a genetically inspired optimization technique that creates a random population of designs in the form of chromosomes, and refines the population for improved performance by calculating the fitness of each design to probabilistically remove weaker design attributes and reinforce stronger design attributes. The chromosomes that represent each design are usually vectors of values, which are modified using genetically inspired operators such as migration, selection, mutation, and crossover to take place between individual genes within each chromosome based on the evaluated fitness of the chromosome. The chromosome representation used in the parametric vehicle design model described in Section 3.1 is setup precisely this way.

GAs are classified as stochastic algorithms, which aid in getting out of local minima in the optimization space. Another stochastic method, simulated annealing, was examined in this experiment, but GAs were found to be more efficient and effective for use in this framework with faster and more reliable convergence. Deterministic methods tend to converge more quickly to the solution, but are more susceptible to getting stuck in local minima. The use of a deterministic method in MATLAB's Optimization Tool Box called fmincon was also examined. As demonstrated with an example in Section 5.4.2.2, this method is highly susceptible to local minima in this design space.

The GA Optimization Tool in MATLAB's Optimization Tool Box was used to generate designs using the Artificial Neural Networks trained on the first stage survey data as the fitness function. Formatting the GA Optimization Tool to breed designs that are compatible with the chromosome vehicle representation described in Section 3.1 requires only to inform the tool that the problem contains 12 variables, and that the upper and lower bounds are 1-by-12 vectors with all values at 100, or all values at 0 respectively. Formatting the ANN, which is trained using MATLAB's Neural Network Fitting Tool, for use with the GA Optimization Tool requires the creation of a proxy function that takes the chromosome in question from the GA Optimization Tool and passes it to a function that retrieves the output from the ANN. The output of the ANN follows the same convention as the survey, where a higher value indicates a higher rating. Because the GA Optimization Tool is looking to minimize this objective function, the proxy function returns the output value of the ANN in cases where the GA is searching for designs with the lowest rating, and the negative of the output value of the ANN in cases where the GA is searching for designs with the highest rating.

5.3.6 DATA ANALYSIS

The survey data from the first stage was trained using MATLAB's Neural Network Fitting Tool. Despite the fact that Levenberg-Marquardt is a deterministic algorithm, the training process for the ANN is stochastic (Prudêncio & Ludermir, 2002). This stochasticity is due to the algorithm using a variety of different random starting weights in an effort not to get hung up in local minima, and due to the tool randomly choosing the subset of data used to train, verify, and test the network. As a result, a good fit is not guaranteed. From the literature (Mitchell, 1997) and looking at similarly dimensioned problems (Tsutsumi & Sasaki, 2008, Hsiao & Huang, 2002) it was suggested that for the complexity of this problem a one hidden layer ANN with between 2 and 20 hidden nodes would be appropriate. Because this research does not intend to advance the art of ANNs, a brute force method was employed to find the optimal network structure and to deal with the stochasticity of training. A search was conducted to find the optimal number of hidden neurons. For each rating category for each participant, eight networks were generated, one with 2, 3, 4, 5, 6, 10, 12, and 20 hidden neurons. The network was trained with a randomly selected 193 (70%) of the 275 samples, and then validated using 41 (15%) samples, and then tested using 41 (15%) samples. The network structure that scored the lowest error for the 41 test samples was assumed to have the optimal number of layers. The optimal number of layers varied from 3 hidden neurons to 20 hidden neurons with no pattern that was discernable to the researcher. To deal with system stochasticity, the network was trained and saved five more times, and the network with the lowest Mean Squared Error for the test data sample was selected to generate verification car designs for the second stage verification survey. Four ANNs were created and trained for each participant, one for each category of judgment.

Once the ANNs were created, the GA tool in MATLAB's Optimization Toolbox was used to find vehicle designs that the ANNs rate as being the highest and lowest in each of the four categories. The GA optimization tool is designed to find the minimum of the fitness function within a set of constraints. The fitness functions used in each case were the four ANNs per participant created to model the survey data from the first stage. In order to find the lowest rated design, the output of the ANN was used as the fitness function, and in order to find the highest rated designs, the negative of the output of the ANN was used as the fitness function. The 20 chromosomes for the initial population of designs used in the GA were randomly generated by MATLAB, and were bred using the default combination of selection, reproduction, mutation, crossover, and migration functions. In order to maximize the chance that the global optimum was found, the default stopping criteria was disabled, and the algorithm was allowed to run for 5000 iterations for each design. The resulting highest and lowest rated designs in each of the four categories for participant 18 are presented in Figures 5.14 – 5.17.

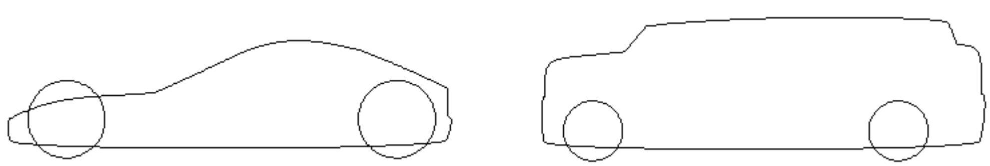

Figure 5.14 – Generated design most sporty (left) and least sporty (right) for participant 18

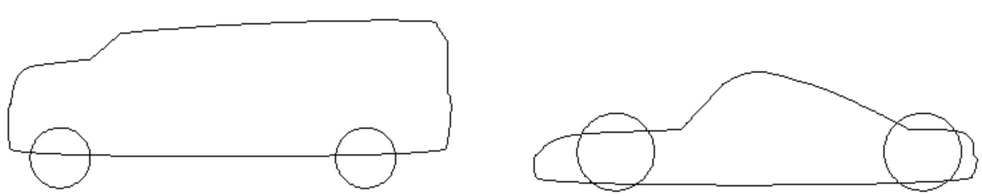

Figure 5.15 – Generated design most rugged (left) and least rugged (right) for participant 18

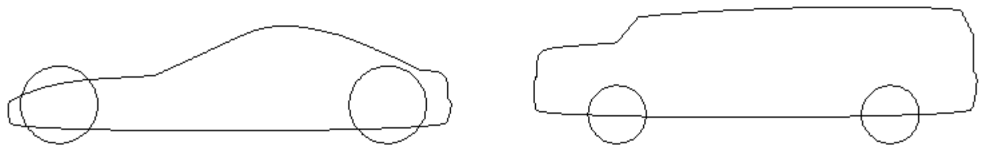

Figure 5.16 – Generated design most beautiful (left) and least beautiful (right) for participant 18

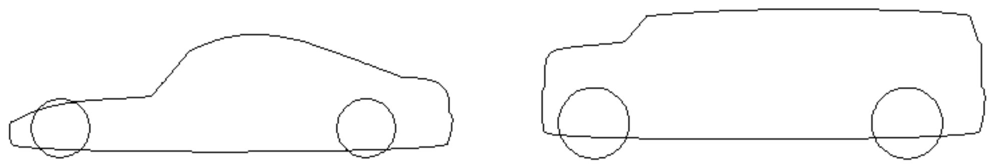

Figure 5.17 – Generated design most fuel efficient (left) and least fuel efficient (right) for participant 18

As described in Section 5.3.3.3, the designs that were generated using the ANN and GA system were assembled into the verification surveys presented in the second stage survey. The generated verification designs were shown three times each to eighteen participants, resulting in 108 total ratings for each category, 54 for highly rated designs, and 54 for low rated designs. These ratings for specifically targeted designs were compared against the average ratings for all survey responses from the first stage survey, for each specific category. These average survey responses contain 276 responses for each participant, or 4968 responses total. Due to the unequal sample size and variance between the 54 specifically targeted design ratings and the 4968 average survey responses, Welch's t-test was used to test for statistical significance. Welch's t-test is a variation from a normal student's t-test that allows for unequal variance and sample sizes between the two populations. Note that the effective degrees of freedom reported here are approximated using the Welch-Satterthwaite equation. For all statistical tests, an alpha level of 0.05 was used (α=0.05).

As described in Section 5.3.3.2, four designs were inserted five times each into the first stage survey, and three times each into the second stage survey. It is hypothesized that participants who are unable to be consistent with ratings in a category across identical vehicle designs presented throughout the study would also not be consistent with the reasoning behind their judgments in that category. Using inconsistent judgments to train the ANN would result in a less accurate model, as well as a less accurate verification survey. The average standard deviation of all random survey data gathered for the study across all categories is 1.04, and the average standard deviation of all ratings targeting specific generated designs is 0.87. It was decided that the average between these two standard deviations, a standard deviation of 0.96, would be used as the cutoff threshold. More specifically, if the average standard deviation between the four cars, presented 8 times each, is above 0.96, the participant would be removed from the study. This filter removed one participant from the sportiness category, three participants from the ruggedness category, nine participants from the fuel efficiency category, and six participants from the beauty category. The data for the study is presented both with and without this filter in the results section.

5.3.7 RESULTS

The average survey ratings from all eighteen participants, without applying the consistency filter, are shown in Figure 5.18. In all four categories, the vehicles generated using the ANN and GA learning and generation system to elicit high and low ratings were rated higher and lower than the average of all 4968 ratings in their respective category by statistically significant amounts. The error bars shown in Figure 5.18 represent the standard error.

Vehicles generated to appear most sporty had significantly higher sportiness ratings than the average ratings in the sportiness category, $t(56.6)=19.49$, $p<0.001$, while vehicles generated to appear least sporty had significantly lower sportiness ratings than the average ratings in the sportiness category, $t(61.6)=24.98$, $p<0.001$.

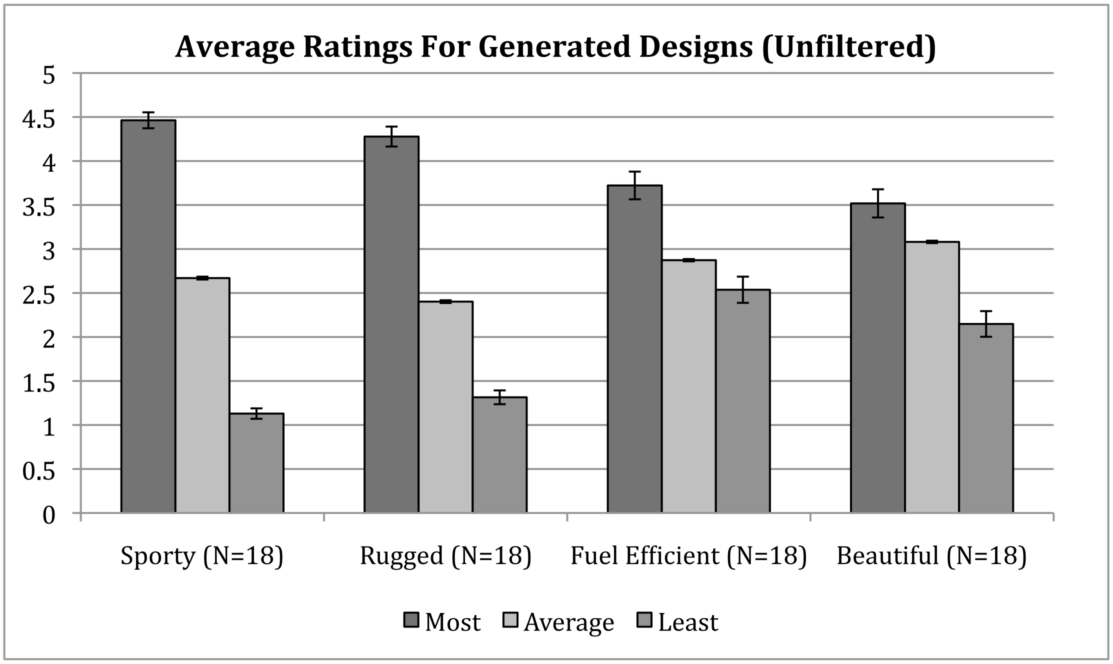

Figure 5.18 – Average ratings for generated designs, unfiltered

Similarly, vehicles generated to appear most rugged had significantly higher ruggedness ratings than the average ratings in the ruggedness category, $t(55.0)=16.38$, $p<0.001$, while vehicles generated to appear least rugged had significantly lower ruggedness ratings than the average ratings in the ruggedness category, $t(57.2)=13.58$, $p<0.001$. Also, vehicles generated to appear most fuel efficient had significantly higher fuel efficiency ratings than the average ratings in the fuel efficiency category, $t(53.7)=5.38$, $p<0.001$, while vehicles generated to appear least fuel efficient had significantly lower fuel efficiency ratings than the average ratings in the fuel efficiency category, $t(53.8)=2.25$, $p=0.014$. Lastly, vehicles generated to appear most beautiful garnered significantly higher beauty ratings than the average ratings in the beauty category, $t(53.8)=2.72$, $p=.004$, while vehicles generated to appear least beautiful had significantly lower beauty ratings than the average ratings in the sportiness category, $t(54.0)=6.37$, $p<0.001$.

A consistency filter was developed to remove participants who were unable to consistently rate the consistency verification designs in the survey. All participants who had a standard deviation greater than 0.96 across the consistency verification designs in each category were filtered out with the assumption that a lack of consistency across the test vehicles would also manifest itself as a lack of consistency with the vehicles used to train the ANN, and a lack of consistency in rating the generated vehicles. The average ratings for the generated designs are shown in Figure 5.19. When compared to the unfiltered results, the differences between the ratings for the generated designs in comparison with the average ratings in each category were even larger. Note that in addition to changes to the averages of the high and low ratings, the means for the average ratings for each category changed as a result of removing the ratings of the filtered participants from the average.

With the filtered results, the results of the significance tests became even stronger. Vehicles generated to appear most sporty had significantly higher sportiness ratings than the average ratings in the sportiness category, $t(53.7)=20.14$, $p<0.001$, while vehicles generated to appear least sporty had significantly lower sportiness ratings than the average ratings in the sportiness category, $t(58.4)=24.57$, $p<0.001$. Similarly, vehicles generated to appear most rugged had significantly higher ruggedness ratings than the average ratings in the ruggedness category, $t(46.5)=19.51$, $p<0.001$, while vehicles generated to appear least rugged had significantly lower ruggedness ratings than the average ratings in the ruggedness category, $t(48.8)=14.59$, $p<0.001$. Also, vehicles generated to appear most fuel efficient had significantly higher fuel efficiency ratings than the average ratings in the fuel efficiency category, $t(27.2)=13.04$, $p<0.001$, while vehicles generated to appear least fuel efficient had significantly lower fuel efficiency ratings than the average ratings in the fuel efficiency category, $t(28.1)=6.19$, $p<0.001$. Lastly, vehicles generated to appear most beautiful garnered significantly higher beauty ratings than the average ratings in the beauty category, $t(36.2)=5.92$, $p<0.001$, while vehicles generated to appear least beautiful had significantly lower beauty ratings than the average ratings in the sportiness category, $t(36.7)=13.20$, $p<0.001$.

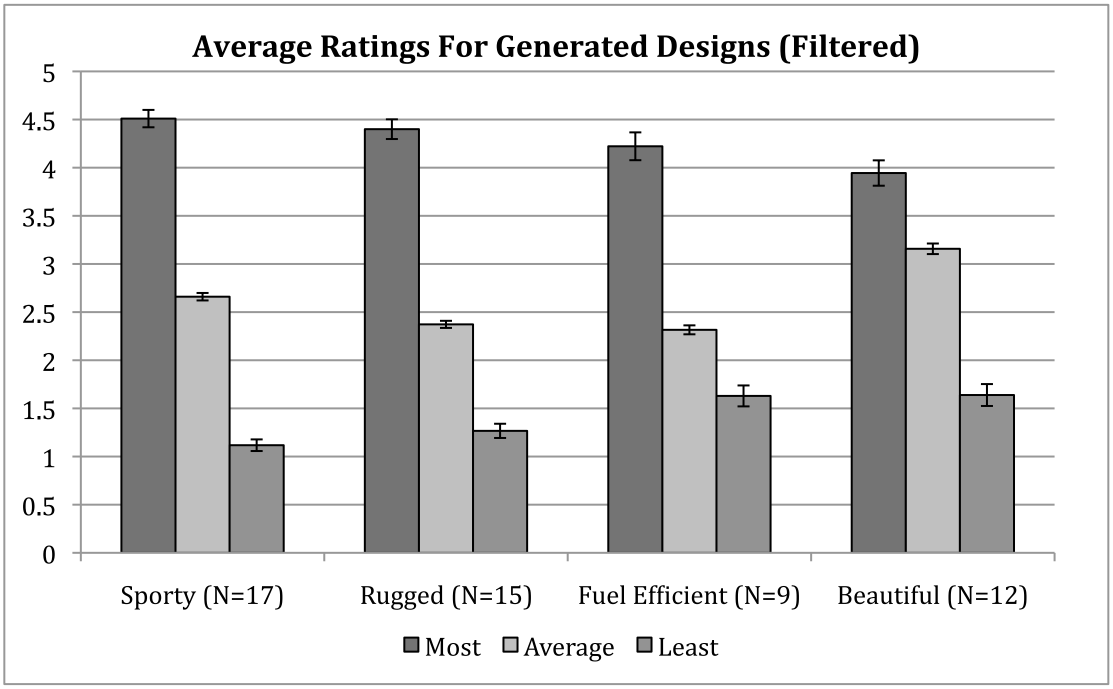

Figure 5.19 – Average ratings for generated designs, with less consistent participants filtered out

5.3.8 DISCUSSION

This study shows that it is possible to train an Artificial Neural Network to mimic human judgments of stylistic form. It was also found that Genetic Algorithms can be used to generate designs that rate high and low from both the ANN and human participants. In all eight cases, designs that were generated with the goal of being rated high or low in each of the four categories did in fact receive ratings that were significantly higher and lower than the average ratings in each category.

In Section 4.1.1.5 an interrater agreement test was applied to the results of that study which found that second to judgments of aerodynamics, which was not tested in this study, participants were most consistent in judging the sportiness of vehicle designs, followed by ruggedness, and were least consistent at judging fuel efficiency. These trends were seen in the results of this experiment as well, with the strongest trends in the data existing for ratings of sportiness, followed by ratings of ruggedness, and somewhat weaker trends for ratings of fuel efficiency.

A consistency verification metric was developed and used to filter the results from this study. This metric removed participants from categories where they were unable to demonstrate sufficient consistency with identical vehicle designs that were presented throughout the surveying process. Applying this filter removed one participant from the sportiness category, three participants from the ruggedness category, nine participants from the fuel efficiency category, and six participants from the beauty category. Removing these participants from these respective categories strengthened the results further by increasing the margin between the targeted high and low ratings and the average non-targeted ratings in each category.

It is clear at this point that the framework developed in this section enables computers to quantify and model human judgments of stylistic form, and to generate new designs to satisfy desires of those human judgments. In addition to the success with targeting the judgments of all participants, it was shown that by removing participants who did not have clear ideas about their judgments of stylistic form, this framework performs with even greater accuracy. This quantification opens the doors for stylistic form to be unified with functional goals in computational optimization, which the next section illustrates.

5.4 TRADEOFFS BETWEEN STYLISTIC FORM AND FUNCTION

This work set out to examine how stylistic form and functional constraints and goals can be brought together into a single cohesive framework. To this end, a computational system was developed that allows an Artificial Neural Network (ANN) to learn human judgments of stylistic form. This network was then coupled with a Genetic Algorithm (GA), which was able to generate new designs that successfully target and elicit specific high and low judgments from human participants in categories of sportiness, ruggedness, fuel efficiency, and beauty. In Section 4.1, a number of tradeoffs between different stylistic form goals and functional goals were identified. This section will demonstrate three of these tradeoffs by computationally generating designs that satisfy these conflicting stylistic form and functional goals.

One of the tradeoffs identified in Section 4.1 was the tradeoff between perceived ruggedness and actual aerodynamic performance. Two other tradeoffs identified in Section 4.1 involved increasing the volume of the vehicle design. One was the tradeoff between increased volume and improved aerodynamic performance, and the other was the tradeoff between increased volume and perceived sportiness.
The first tradeoff is between two functional goals that interact via the form of the vehicle design, while the latter tradeoff is between a stylistic form goal and a functional goal.

In the analysis of tradeoffs between multiple objectives, there may not be one specific optimal design, but rather a family of optimal designs that represent different choices and compromises in the tradeoffs between different objectives. Having a way to visualize or map optimal designs along the tradeoff can be useful, especially in cases where specific goals for the objectives are unknown or have not been finalized. The optimal set of solutions wherein no improvements can be made to any objectives without hurting the performance of another objective is called the Pareto Set or Pareto Frontier (Fudenberg, 1983; Fonseca, 1995). A plot of this Pareto Set can help to visualize the tradeoffs between multiple objectives.

5.4.1 TRADEOFFS BETWEEN RUGGEDNESS AND AERODYNAMIC PERFORMANCE

In this section, the tradeoff between perceived ruggedness and aerodynamic performance in highly rugged cars is analyzed with the help of a plot of the Pareto Set. The Pareto Set plotted in Figure 5.20 was found using the Multi-objective Genetic Algorithm Tool in MATLAB's Optimization Toolbox. The purpose is to study the tradeoff in cars that are perceived to be highly rugged, so a subset of the data is plotted from ruggedness ratings of 3.8 and higher. In order to plot the tradeoffs between two objectives, two separate objective functions are needed. To represent judgments of perceived ruggedness, the Artificial Neural Network trained on survey data from participant 15 in Section 5.3 is used as the first objective function. Since the Multi-objective Genetic Algorithm Tool seeks to minimize all objective functions, the negative of the output of the network was used. Note that the absolute value of the ruggedness rating objective function was plotted for clarity. In order to represent aerodynamic performance, the parametric aerodynamic model described in Section 3.2 is used as the second objective function. All points on this plot refer to an optimal design based on a chosen tradeoff between ruggedness and aerodynamic performance.

It is interesting to notice that due to the characteristics of this specific Artificial Neural Network, the plotted Pareto Set appears to be a piecewise response. A closer look at the vehicle designs at each of the four marked points suggests why. Starting from the upper right of Figure 5.20 with the highest achieved level of ruggedness rating of 6.04 and an aerodynamic coefficient of drag of 0.320 at data point A, the corresponding vehicle design for data point A is shown in Figure 5.21(a). In traversing from data point A to data point B, there appears to be a smooth transition trading some ruggedness for gains in aerodynamic performance, with diminishing returns of how much aerodynamic performance is gained per unit of ruggedness lost when approaching data point B. The ruggedness for the design at data point B has been reduced to 4.68, while aerodynamic performance has improved to 0.303. The vehicle design at data point B is shown in Figure 5.21(b).

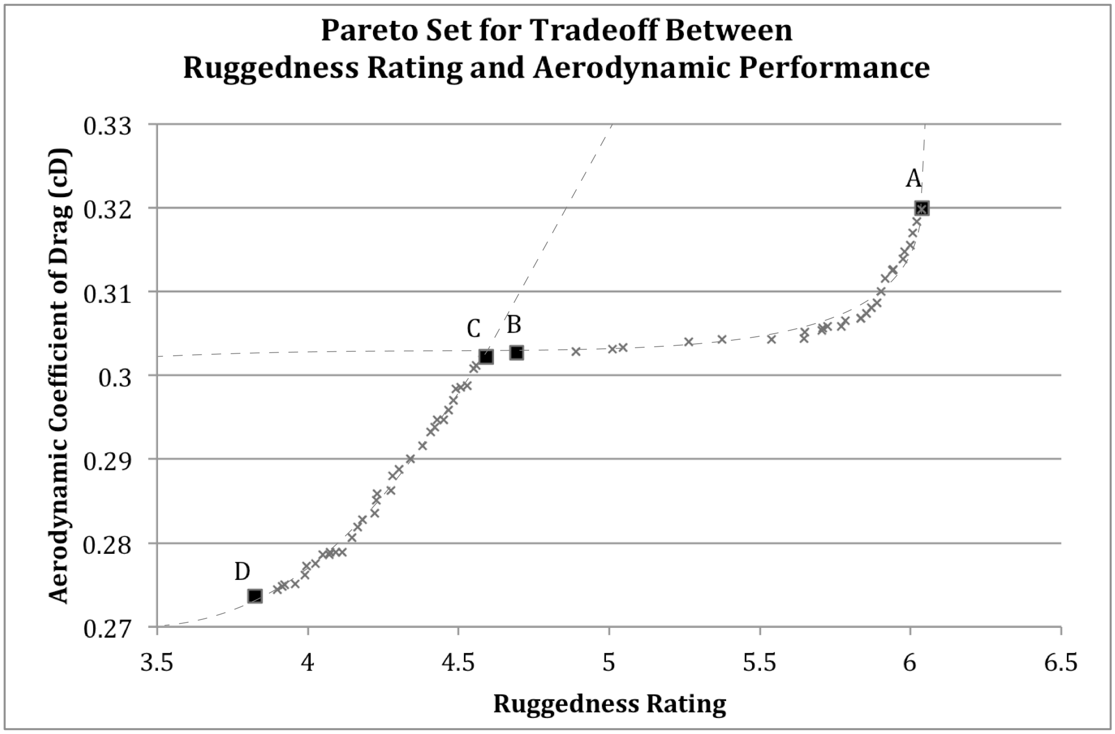

Figure 5.20 – The Pareto Frontier of optimal tradeoffs between the ruggedness rating and aerodynamic performance

In studying the vehicle designs (not shown) between data point A and data point B a smooth and continuous transition between the two designs is seen, with the majority of the change happening as a result of a reduction of body height. These designs belong to a single design family that falls along the dotted line traveling through data points A and B in Figure 5.20. At data point C a discontinuity in slope is observed. The vehicle design at data point C is shown in Figure 5.21(c) and shows a discontinuous jump from the vehicle design progression observed between data point A and data point B. The rear windscreen at data point C becomes much more highly raked, and the body height has increased above that of data point B. A closer study of the models indicates that this increased raking of the rear windscreen causes a reduction in perceived ruggedness, but also an increase in aerodynamic performance, resulting in almost exactly the same ruggedness rating and aerodynamic performance as the vehicle design at data point B. Where these the two design families (depicted by the two dotted lines) intersect, where are two different but equivalent vehicle designs that have the same ruggedness rating and aerodynamic performance, belonging to each of the two design families. At ruggedness ratings above 4.6, designs in the first design family, which contains data points A and B are more optimal, whereas below ruggedness ratings of 4.6, designs in the second design family, which contains data points C and D are more optimal. The Pareto Set is made up of designs from both design families, but only from regions where the respective design family is optimal. The design at data point C has a ruggedness rating of 4.59 and an aerodynamic coefficient of 0.302. In studying the vehicle designs between data point C and data point D (not shown), which both belong to the second design family, another smooth and continuous transition between the two cars is seen, with the majority of the change once again happening as a result of a reduction of body height, resulting in the vehicle design at data point D, which is shown in Figure 5.21(d). Once again there is a relationship where perceived ruggedness is being traded for an improvement in aerodynamic performance. In the tradeoff shown between data point C and data point D, the effect of diminishing returns appears not to be as strong as the one seen between data point A and data point B. The perceived ruggedness at data point D is 3.81, and the aerodynamic coefficient is 0.274. This mapping of the Pareto Set allows designers to better visualize the design options that are afforded to them, and can help designers to make better design decisions, especially pertaining to tradeoffs and synergies between design goals, and identifying different families of designs.

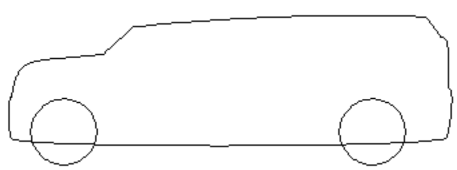

Figure 5.21(a) – Pareto optimal solution (a)

Figure 5.21(b) – Pareto optimal solution (b)

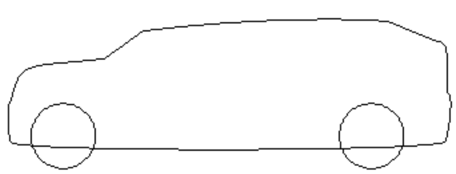

Figure 5.21(c) – Pareto optimal solution (c)

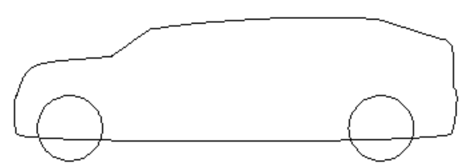

Figure 5.21(d) – Pareto optimal solution (d)

5.4.2 TRADEOFFS BETWEEN VEHICLE VOLUME AND AERODYNAMIC PERFORMANCE

In the current climate of high fuel prices and people's desires for capacious and useful vehicles, it should come as no surprise that tradeoffs between a vehicle's interior volume and a vehicle's aerodynamic performance are of utmost concern. In the previous section a Pareto Set was generated to map the interface between two tradeoffs. While this is a wonderful way for a designer to visualize tradeoff possibilities when the desired goals are unknown, the method can be computationally expensive when done on the computer, or resource intensive when done manually. In some design problems specific target goals are known, while others may not be. This section will focus on designing a series of vehicles with specific goals using the Artificial Neural Network and Genetic Algorithm based generation tools developed in Section 5.3. This section demonstrates how this tool can be manually used to target specific design goals.

In general, GAs require fitness functions that provide a single scalar value for a given set of parameters. This can pose difficulties in multi-objective optimization, where the different objectives can be represented in substantially different ways. There have been several approaches to combine multiple objectives into a single scalar function that is appropriate for use in GAs (Fonseca, 1995; Zitzler, 1999; Coello, 2000). Weighted sum approaches are one common approach where multiple objectives are adapted into a single fitness function made up of a sum of normalized positive scalar objective functions, which can then be weighted if desired (Zadeh, 1963; Goicoechea et al., 1982). Another method is goal programming (Ignizio, 1976; Schniederjans, 1995) where one objective is minimized while the remaining objectives are constrained at goal values. This process switches between objectives using a variety of goal programs. Another method is target vector optimization (Hans, 1988), which requires a vector of goal values. In that case, optimization is driven towards the shortest squared sum vector difference between any candidate solution and the goal vector.

5.4.2.1 Optimizing Aerodynamic Performance with Fixed Volume

The designs generated here focus on two purely functional goals: aerodynamic performance and vehicle volume. More specifically, the goal of this design problem is to generate the most aerodynamic shapes that can be achieved to satisfy a given volume. Because the volume is constrained in this case, and the aerodynamic coefficient of drag is to be minimized, the fitness function must be set up accordingly. The fitness function for the GA in this problem is set up using a weighted sum approach (Zadeh, 1963; Goicoechea et al., 1982). When using a weighted sum approach, the fitness function is comprised of multiple objective functions that are normalized to span a common range of difference between maximum and minimum values, and manual weights for the objective function that can be added to adjust the preferences of the system (Marler & Arora, 2004). It is important to note that the actual values of the objective functions are unimportant; rather it is the difference between the maximum and minimum that is spanned by the objective function that is normalized. The first objective function in this problem, the aerodynamic coefficient of drag, which normally has a function range of approximately 0.20 to 0.50 was multiplied by 333 to be normalized to a range of 66 to 166, a difference of 100. The second objective function is structured differently from the first because the volumetric assessment in this model has a target value. Because of this target value, the objective function should no longer penalize the fitness function after its value drops below the desired target. In this case the objective function is the negative of the volumetric assessment function because the goal is to maximize the volume, and a Boolean operator that sets the volumetric assessment function equal to the desired volume when the desired volume is reached, removing the effect of this objective function. The range of this function was also normalized for a difference between maximum and minimum values of approximately 100.

Using this new fitness function, the GA generated the design pictured in Figure 5.22(a) as the most aerodynamic shape that can be achieved with a minimum volume of 65 cubic feet with a coefficient of drag of 0.196. By increasing the targeted minimum volume to 75 cubic feet, the design pictured in Figure 5.22(b) was designed as the most aerodynamic shape with a coefficient of drag of 0.223. Increasing the targeted minimum volume one last time to 85 cubic feet results in the design pictured in Figure 5.22(c) as the most aerodynamic shape with a coefficient of drag of 0.254. Using the methods discussed here, a GA in conjunction with two functional design models is able to generate functional designs that target specific performance targets. One interesting note is that the most aerodynamic car design that can encompass a larger interior volume appears to resemble the design shown in Figure 3.15, which was designed to resemble a Toyota Prius, a real vehicle that was designed to balance goals of aerodynamic performance and spaciousness.

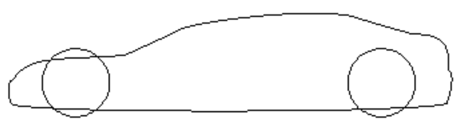

Figure 5.22(a) – Most aerodynamic, 65 cu.ft.

Figure 5.22(b) – Most aerodynamic, 75 cu.ft.

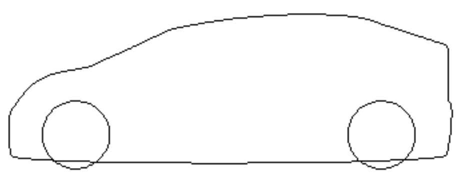

Figure 5.22(c) – Most aerodynamic, 85 cu.ft.

5.4.2.2 Deterministic Search and Local Minima

The optimization process using Genetic Algorithms to generate vehicle designs in this framework is slow and computationally expensive. It is therefore tempting to consider the use of a deterministic search method such as MATLAB's fmincon function, which is faster and less computationally expensive than GAs. The drawback to deterministic search methods is that they are more susceptible to getting stuck in local minima. One such example is shown here. Fmincon was used to solve the problem in the previous section to design the most aerodynamic shape that encompasses the volume of 85 cubic feet. Using the GA in the previous section to solve this problem, a vehicle with a coefficient of drag of 0.254 and a volume of 85 cubic feet was generated as shown in Figure 5.22(c). Fmincon requires a starting point to begin its search. With a chromosome with all parameters equal to 50 as the starting point, the algorithm found the local minimum pictured in Figure 5.23, which has a volume of 83.25 and an aerodynamic drag coefficient of 0.268, which is inferior on both accounts to the global solution found by the GA in the previous section. A look at the chromosome of the vehicle reveals that the local minimum that was found had the first nine parameters, which are used in the shaping of the body, at the maximum and minimum parameters values. Randomly starting with other initial parameter values did eventually yield the optimal solution found using the GA in the previous section.

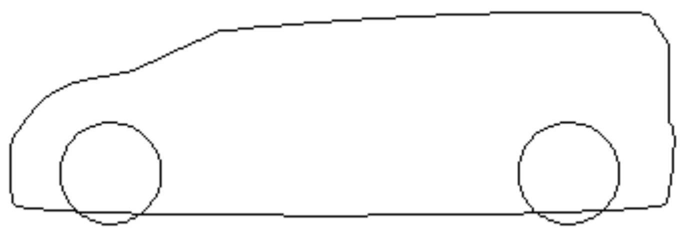

[100 100 0 100 0 0 0 100 0 50 50 50]

Figure 5.23 – Local minimum when using fmincon

5.4.2.3 Adding Sportiness to a Useful Design

In Section 5.4.2.1, three vehicles were designed to illustrate the powerful computational design abilities that can be unveiled by coupling a GA with two functional assessment models. The vehicles designed were of a purely functional shape. The goal of this dissertation is to couple stylistic form and function in a unified approach that allows for synergies and optimal tradeoffs. Using the methods described in Section 5.3, this section aims to computationally add some sportiness to the three designs generated in Section 5.4.2.1.

The ANNs that were trained in Section 5.3 differed greatly between participants. In order to illustrate the range of differences modeled, this section will demonstrate how the GA generates vehicles that add sportiness to the previous set of goals using the ANNs from two different participants for comparison.

In Section 5.4.2.1, an aerodynamic coefficient of drag of 0.254 was achieved for a vehicle with 85 cubic feet of volume. This presumably is the theoretical maximum in this model for the best coefficient of drag a vehicle with that volume can achieve. In this design problem, the goal is to design the sportiest designs that balance a targeted vehicle volume with an optimal aerodynamic performance that also must perform at least as well aerodynamically (cD = 0.254) as the 85 cubic foot vehicle designed in Section 5.4.2.1.

In this problem, three different types of objective functions are used. In the first objective function, the sportiness rating, is maximized. For compatibility with MATLAB's GA Optimization Tool, which aims to minimize objective functions, the objective function must be formatted such that the negative of the sportiness rating is minimized. The sportiness rating is assessed through the ANN. The second objective function is to maximize the volume of the vehicle until specific target values have been achieved. Much like in Section 5.4.2.1, the negative of the objective function is to be minimized until the target value is reached, at which point a Boolean operator enforces the target value. The third objective function is the aerodynamic coefficient of drag, which is to be minimized, and must not exceed 0.254. This maximum coefficient of drag is enforced using a penalty function (Marler & Arora, 2009), which adds a large penalty value to the objective function value if the coefficient of drag exceeds 0.254. All objective functions were normalized to a range between maximum and minimum values of 100 and were summed together without additional weighting. The ANNs used to model sportiness in this problem are from Participant 3 and Participant 18. The sportiest design for Participant 18 without any other objective functions is shown in Figure 5.14. Note that because the ANNs for the two participants are different and have different ranges, the normalization weights used were also different in order to achieve the desired range of 100.

First, two vehicles were designed for a target volume of 65 cubic feet. The design based on the sportiness ANN from Participant 3 is shown in Figure 5.24(a). This design achieves the target volume of 65 cubic feet, and manages it with a sportiness rating of 7.03, and an aerodynamic performance of 0.211. Similarly, the design based on the sportiness ANN from Participant 18 is shown in Figure 5.24(b). This design also achieves the target volume of 65 cubic feet, and manages it with a sportiness rating of 5.31 and an aerodynamic performance of 0.202.

Next, two vehicles were designed for a target volume of 75 cubic feet. The vehicle designed using the sportiness ANN from Participant 3 is shown in Figure 5.24(c).

This vehicle achieves the target volume of 75 cubic feet, and manages it with a sportiness rating of 6.86 and an aerodynamic performance of 0.249. The vehicle designed using the sportiness ANN from Participant 18 is shown in Figure 5.24(d). This vehicle achieves the target volume of 75 cubic feet, and manages it with a sportiness rating of 4.84 and an aerodynamic performance of 0.239.

Lastly, two vehicles were designed for a target volume of 85 cubic feet. The vehicle designed using the sportiness ANN from Participant 3 is shown in Figure 5.24(e). Due to the restrictions on aerodynamic performance and sportiness, this vehicle is not able to reach its target volume of 85 cubic feet; instead it measures 79.96 cubic feet. This increase in volume has also reduced the sportiness rating to 5.37 and the aerodynamic performance is at the limit of the penalty function at 0.254. The vehicle designed using the sportiness ANN from Participant 18 is shown in Figure 5.24(f). This vehicle also was unable to achieve the volumetric goal of 85 cubic feet, instead achieving only 82.59 cubic feet, while the sportiness rating has dropped to 4.45, and the aerodynamic performance has also been pushed to the limit of the penalty function at 0.254.

In this section six vehicles were designed that balance the functional goals used in Section 5.4.2.1 with the desire to build sporty vehicles based on learned preferences through the ANN. ANNs from two different participants were used to generate the vehicles in this section. As a result, two strikingly different sets of vehicle designs resulted based on judgments of what each participant feels is the sportiest design. While the introduction of the sportiness objective prevented the designed vehicles from achieving the 85 cubic feet volumetric target, two uniquely and differently designed vehicles that show optimal tradeoffs with the three goals were generated.

The experiment in this section also helps to confirm the fact that the ANN does in fact capture interaction effects in stylistic form, a benefit that the logit model, as implemented in Section 5.1, could not. This is illustrated by studying the position and size of the wheels in each vehicle design. Wheel size and wheel position are not active variables in either the aerodynamic model or the volumetric assessment. As the changes in volume and aerodynamics were enacted on the vehicle designs resulting in shape changes, the sportiest size and positioning of the wheels changed accordingly. If interaction effects were not modeled, the wheels would remain stationary throughout all body shape changes for cars built with each participant's ANN. In checking for the effects of the stochasticity of the GA used, each vehicle was designed multiple times using the GA, and the shift in size and position of the wheels were consistent throughout all trials, showing that the changes were in fact caused by modeled interactions effects in the ANN, and not the stochasticity of the GA.

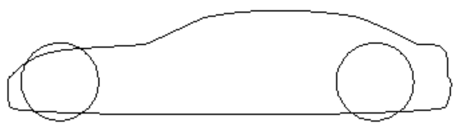

Figure 5.24(a) – Participant 3 65 cu.ft.

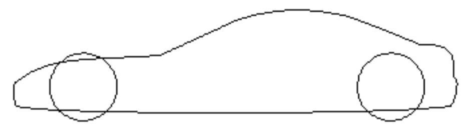

Figure 5.24(b) – Participant 18 65 cu.ft.

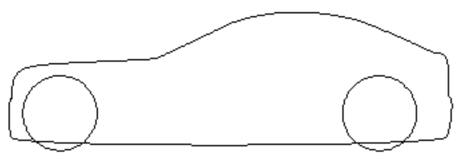

Figure 5.24(c) – Participant 3 75 cu.ft.

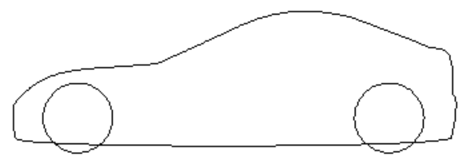

Figure 5.24(d) – Participant 18 75 cu.ft.

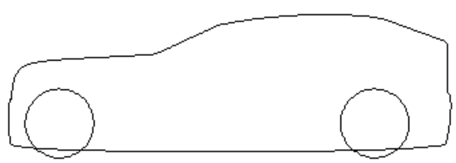

Figure 5.24(e) – Participant 3 79.96 cu.ft.

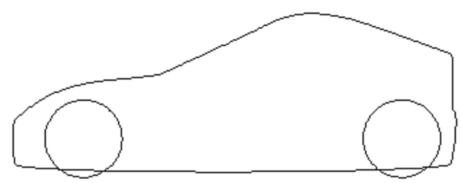

Figure 5.24(f) – Participant 18 82.59 cu.ft.

CHAPTER SIX

UNDERSTANDING AND INSPIRING DESIGNERS

How do designers represent and tradeoff between multiple goals in a form and function design problem? Can presenting the problem a different way, or giving additional information inspire a designer or a group of designers to interface between these goals more effectively? And under what circumstances are designers most effective at incorporating this new information into problem solving? Often times in solving problems that incorporate multiple objectives, such as stylistic form and function, designers need to be able to apply new knowledge about the problem or about relationships between different objectives in the problem, which can be highly beneficial to the effective management of tradeoffs or exploitation of synergies between the multiple objectives. Unfortunately, designers may not always successfully make a connection between this new knowledge and the problem they are solving (Forbus, Gentner, & Law, 1995; Gick & Holyoak, 1980). Furthermore, as is shown in Section 6.2, even if a connection is made, the knowledge may not be well applied for maximum benefit. In order to better understand how designers assimilate and apply newly acquired information in a design problem, and how this newly acquired information affects problem solving, a cognitive study was conducted and is described in Section 6.1. The findings of this experiment then inspired the experiment described in Section 6.2, which studies how example images can help designers to more effectively apply representation and goal knowledge to enrich the solving of a form and function design problem.

6.1 THE EFFECTS OF OPEN GOALS AND ANALOGICAL SIMILARITY ON DESIGN PROBLEM SOLVING

An experiment was conducted to gain an understanding of how people assimilate and apply newly acquired information when ideating solutions to a design problem by studying how the timing and nature of problem-relevant information can affect idea generation in an open-ended design problem. More specifically, the effects of presenting surface similar information before design conceptualization, or surface dissimilar information before and during design conceptualization on the quantity, breadth, and novelty of solutions generated were analyzed. The effects of open goals, fixation, and priming, as well as their implications in design problem solving are examined. It was found that information that is more distantly related to the design problem impacted idea generation more when there was an open goal to solve the problem, while information that is more obviously similar to the problem impacted idea generation more than distantly related information when seen before problem solving has begun. Evidence of induced design fixation and priming were also observed. For more information about this study, please see Tseng et al (2008).

We have observed professional designers breaking from the conceptualization process at points of frustration or impasse to browse magazines or surf the web, seemingly with no specific purpose. When returning to the ideation process, new concepts begin to emerge. This research contributes to the literature of foundational cognitive principles that inform the design process. In particular this work studies the types of analogies that most impact design creativity and the timing when seeking such analogical stimulation is more effective.

It has been shown that designers are particularly susceptible to information from example solutions such as existing products that are similar to what is being designed (Chrysikou & Weisberg, 2005; Jansson & Smith, 1991; Perttula & Liikkanen, 2006; Purcell & Gero, 1996). Designers have even been observed to incorporate poor aspects of existing solutions into their own solution (Jansson & Smith, 1991). One possible explanation for this is that designers become fixated on existing design solutions to the extent that they are not able to think of any other ways to solve the current problem. In this situation, fixation on existing solutions could prevent the designer from being able to come up with an innovative solution to the problem. While these findings may be useful in routine design when similar products already exist, new design problems seldom come with example solutions. Instead, designers often subconsciously look to other devices that they have encountered or may encounter while working on the problem.

Some theories of creativity posit that the source of creative ideas is the combination of distantly related concepts (Campbell, 1960; Simonton, 1999). Perhaps if designers were able to think of distant but relevant ideas, they could avoid becoming fixated on existing solutions. However, research has shown that people are not very good at retrieving and using information that is analogically related to the problem they are trying to solve (e.g., Forbus, Gentner, & Law, 1995; Gick & Holyoak, 1980). These findings lead to the conclusion that people only rarely make use of distantly related information when they are trying to solve a problem.

However, it has been noted that much of this work on analogical transfer has made use of an experimental design where people learn about some material and then later attempt to solve a problem where the learned material could be analogically mapped on to the problem to help solve it. Alternatively, people could encounter relevant information during a break in problem solving that may lead to a higher rate of analogical mapping (Christensen & Schunn, 2005). In fact, it has been shown that having an open goal to solve a problem, a goal which has been set but one for which the associated task has not been completed, leads to the implicit acquisition of relevant information even while not working on a problem (Moss, Kotovsky, & Cagan, 2007). Additionally, people may be most sensitive to new information around the time when they reach an impasse on a problem (Moss, Kotovsky, & Cagan, 2008).

In research on analogy, a distinction is often made between surface similarity and structural or deep similarity (e.g., Forbus et al., 1995; Holyoak & Koh, 1987). Surface similarity is similarity in appearance or attributes. For example, a bicycle may bear some resemblance to a pair of glasses when viewed from the side or two math word problems may both involve similar objects like apples and oranges. Structural similarity, however, means that two things involve similar relationships. For example, the atom and the solar system involve a similar configuration of objects, but they are not similar in appearance. Two math problems may be similar on the surface as noted, but when one involves calculating the total amount of fruit and the other involves calculating the probability of picking an apple out of a bin of apples and oranges, then the two problems are structurally different.

In design, devices can be similar in appearance, purpose, or function. Here, purpose is defined as the main way in which the device is used while function involves a more abstract view of what the device is doing. Two different types of clocks may be highly similar in function, purpose, and appearance. A clock and a watch may be similar in function and purpose but less similar in appearance. However, a bathroom scale and a pressure gauge may be similar in function (i.e., measuring a force, or force per unit area), but not at all similar in appearance or purpose.

Based on the results in the analogy literature described above, designers may find it difficult to recognize analogically useful information from past design experiences if the relationships between the experiences and problem bear structural similarity (i.e., functional similarity) but little or no surface similarity (i.e., appearance or purpose). In the case where the problem solver has the goal to solve a problem but has not yet completed the solution, the problem solver has an open problem-solving goal. Since having an open goal makes it more likely that relevant information is incorporated into problem solving even when the person is not actively engaged in solving the problem, designers may be better able to make the connection between this same information and the problem if they see the information after problem solving has begun.

It has also been found that general representations of analogous information are more likely to be applied to cross-domain design problems than domain specific representations (Linsey et al., 2007). For instance, framing an air mattress as "a device that uses a substance from the environment it is used in", rather than "a device that is filled with air" makes it more likely to be used later in relevant design problems. So, encountering information that leads to a more general framing or representation of the information may make it more likely to be used while solving a design problem.

One of the main goals of this experiment was to examine whether people are able to better recognize and use relevant principles from sources that are not obviously related to the problem (i.e., items that share functional characteristics but not purpose or appearance) when they have an open goal. To examine this, surface dissimilar information that was structurally similar was presented to problem solvers (designers) either before conceptualization (problem solving) or during a break in conceptualization. In addition, this surface dissimilar information was presented as a group of different devices to encourage a more general representation of the information. Another goal was to assess whether principles from surface similar sources presented before problem solving affect problem solving more than from surface dissimilar sources, and so a condition where surface similar information was presented before problem solving was compared to the case where surface dissimilar information was presented before problem solving.

6.1.1 HYPOTHESIS

Three hypotheses were examined in this experiment: 1) devices which are more distantly related to the problem would impact idea generation more when there was an open goal to solve the problem, 2) information which is more obviously similar to the problem would impact idea generation more than distantly related information when seen before problem solving has begun, and 3) functional principles of the presented designs would appear more frequently in the solutions of the participants who saw those designs than in those participants in the control condition who received no problem-relevant material.

6.1.2 METHODS

The problem used in this experiment was an open-ended design problem where participants were asked to generate conceptual designs for as many time-keeping devices as possible using only a provided list of household objects. There are two key comparisons for our hypotheses: 1) comparing highly related and distantly related information before problem solving has begun and 2) comparing distantly related information given before problem solving to when the same information is given during a break in problem solving. Three conditions were designed which allowed us to assess these comparisons, and in these conditions we manipulated the timing of when problem-relevant information is given (before problem solving or during a break in problem solving) and whether the presented problem relevant information contains surface similarities or structural similarities. In addition, a control condition was included as a baseline in which participants only saw irrelevant information. The problem relevant information that was presented was one of two sets of device descriptions. One set consisted of a description of three clocks, and this set was highly similar to the presented problem in function, purpose, and possibly appearance. The other set consisted of descriptions of three distant devices that were not similar to the design problem in appearance or purpose, but in which some of the functional information could be used to solve the design problem.

6.1.2.1 Participants

Seventy-one Carnegie Mellon University undergraduate seniors in the Department of Mechanical Engineering were recruited from two senior courses and voluntarily participated in this experiment.

6.1.2.2 Design and Materials

All participants solved the same design idea generation problem, which is shown in Figure 6.1. The timing and type of problem relevant information given to the participants was manipulated. There were two times when information was presented: before the problem solving began, labeled "pre-problem", and during a break five minutes after problem solving began, labeled "during-break". Each participant was presented with information at these times. The information could either be irrelevant to the problem, a description of three clocks (the surface similar information), or a description of three distant devices (the surface dissimilar information). The three distant devices used were a water meter, a heart rate monitor, and a cassette tape recorder. The irrelevant information, or filler task, consisted of three short summaries of current news stories. The device descriptions are shown in Appendix B. The design of all four conditions is shown in Figure 6.2. Participants were randomly assigned to one of these four conditions.

Participants in the control condition (N=18) were presented with the filler task for both the pre-problem and during-break reading tasks. In the clocks-before condition (N=17) participants were given the clock descriptions for the pre-problem reading task, and the filler task for the during-break reading task. In the devices-before condition (N=18) participants were given the descriptions of distant devices for the pre-problem reading task, and the filler task for the during-break reading task. In the devices-during condition (N=18) participants were given the filler task for the pre-problem reading task, and the descriptions of the three distant devices for the during-break reading task. The filler task used for the clocks-before condition, devices-before condition, devices-during condition, and for one of the control condition breaks was the same.

6.1.2.3 Procedure

The experiment was run in groups in two consecutive class times, with 41 participants in the first class and 30 in the second. Participants received visibly identical packets that contained all materials. Each task was contained in a separate envelope within the packet labeled A, B, C, and D to be used in sequence. The participants were verbally instructed between tasks to advance from envelope to envelope, and were only allowed to view the materials in the current envelope at any one time.

The clock is one of the oldest human inventions, requiring a physical process that will proceed at a known rate and a way to gauge how long that process has run. As the seasons and the phases of the moon can be used to measure the passage of longer periods of time, shorter processes had to be used to measure off hours, minutes, and seconds.

You need to come up with as many of these shorter processes to measure the passage of hours, minutes, and seconds as you can in ten minutes. The time measurement does not have to be in any known unit so long as it is repeatable so that you can repeat it with a clock at a later time.

You are alone in a large featureless room with no windows, a door with doorknob, a hanging light fixture on the 10-foot ceiling, and a sink and drain with working tap.

The only other items in the room are:

Three rolls of adhesive tape	a blue click-type ballpoint pen
a roll of twine	a 12” wooden ruler
a 1 qt Tupperware container with lid	a 3 kg lead weight with hook
a gallon metal can of black latex paint with lid	a 8” tall candlestick with holder
a 2” wide paint brush with wooden handle	a box of matches
a 7 foot aluminum ladder	a thermometer
a 6” serrated hunting knife	a handle (large bottle)of vodka

Please draw or describe the concept of your solutions in order in the boxes provided and mark the time as projected by the laptop in the front of the classroom to the second (hh:mm:ss) in the space provided when you finish each solution. More pages are attached as needed.

Figure 6.1 – Idea generation problem

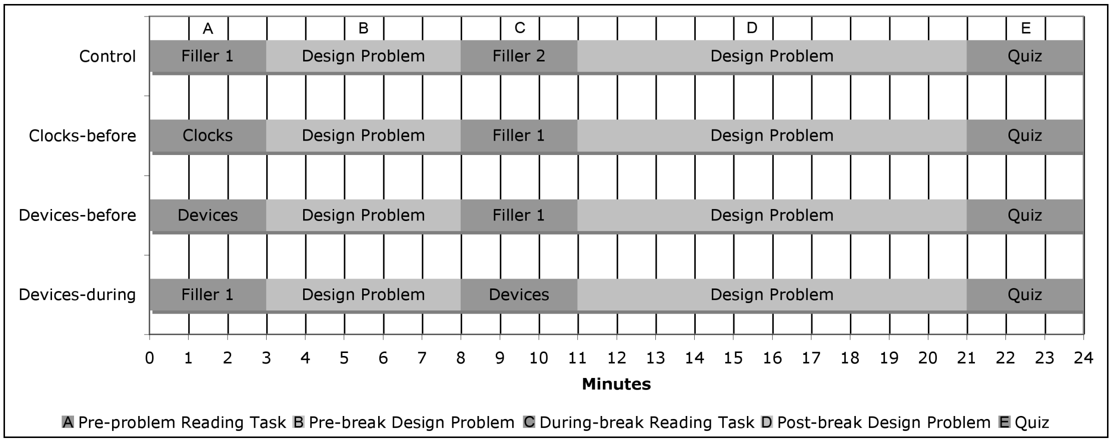

Figure 6.2 – Experiment design for all four conditions

Each participant began with the three-minute reading task, which was specific to his or her randomly assigned condition. Next, all participants were given five minutes to work on the design problem. All participants were instructed to draw or describe their solutions consecutively in the boxes provided and to label each box with the time they finished the solution in hh:mm:ss format, as projected in the front of the classroom. The format with two sample solutions can be seen in Figure 6.3. Fourteen boxes were provided for each problem solving session, and no participant reached this limit. The participants were encouraged to generate as many solutions as possible. After the five minutes, the participants were given a break from problem solving during which they were given three minutes for the second reading task. After the break, all participants were given an additional ten minutes to continue work on the design problem in the same format as before. The participants were verbally instructed not to write down the same answers as before but told that these solutions should be in addition to the previous solutions from the first 5 minutes. The participants were not allowed to look back at their previous solutions. At the end, all participants were given a previously announced quiz to assess whether they retained the information from the two reading tasks to ensure that they read the material and that any failure to use the material in problem solving was not due to an inability to remember the information.

6.1.2.4 Data Analysis

All solutions were analyzed using the participants' drawings and descriptions. Each solution was categorized according to the function(s) used to tell time in the design. This categorization resulted in fifteen functional categories that were found to fit 97%

of all solutions generated. For example, the solution shown in Figure 6.3(a) was categorized as a “Rate of Heating/Cooling” solution. The remaining three percent were lumped into a sixteenth category of “other” solutions. Solutions that included principles from multiple functional categories were placed in all relevant categories in fraction (4% of the solutions). For instance, a solution that uses the sink to fill a container in a see-saw arrangement to offset the 3 kg weight on the other side, as seen in Figure 6.3(b), would be placed half in the “rate of flow/fill” category and half in the “weight equilibrium” category. Invalid solutions were defined as ones where the design used a component not given in the problem statement, or where the description of the device was incomplete or abandoned. These invalid solutions, which made up less than 4% of the data, were excluded from analysis. The resulting average number of designs per participant in each category can be found in Table 6.1. All data was first coded by one researcher as described above, and then the designs generated by five participants from each condition (28% of the data) were randomly selected and were coded by another researcher. The two researchers showed 89% agreement and a Cohen’s Kappa of 0.87, which supports the use of this coding system as a reliable way to categorize the data.

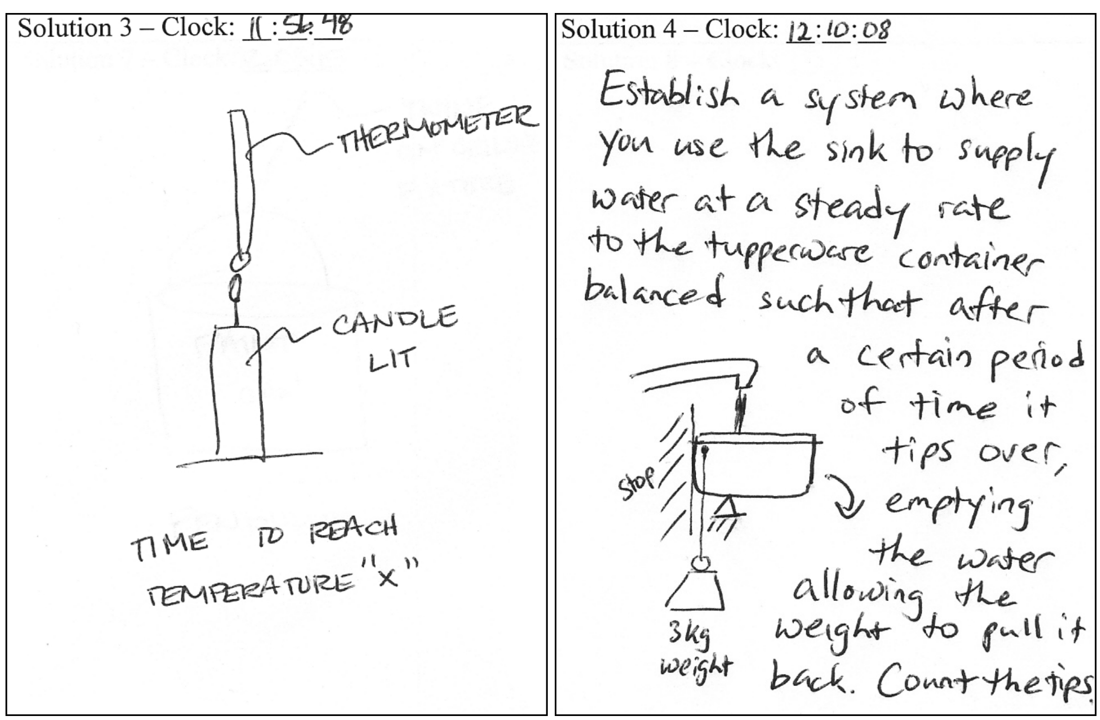

(a) **(b)**

Figure 6.3 – Example solutions **(a)** rate of heating/cooling solution and **(b)** multi-category solution

Using this categorization, four dependent measures were defined. 1) The *total number of designs* is the number of solutions generated by each participant in both the five minute pre-break and ten-minute post-break time periods. 2) The number of *functional repeats* is the number of number of times a participant generated a solution in a functional category in which they had already generated a solution. Solutions that spanned multiple categories were only counted as a repeated design if both solutions were categorized in exactly the same set of categories. 3) The number of *functionally distinct* designs is the number of different categories a particular subject generated at least one design in. Note that the sum of a participant's functionally distinct designs and repeated designs is equal to the total number of designs generated by that participant. 4) The *novelty* of each solution is a measure of its uniqueness across all participants' solutions and was measured by adapting an originality metric defined by Jansson and Smith (1991). The novelty of a particular design is found as the sum of the 'n' scores for an individual's ideas divided by the number of ideas generated for that participant. The 'n' score for each item was calculated across all conditions as:

$$n = 1 - \frac{\text{number of functionally similar designs generated by other subjects}}{\text{total number of designs for all subjects}}$$

Two designs were considered functionally similar designs if they were both assigned to the same functional category.

6.1.3 RESULTS

The average total number of designs, number of functional repeats, and the number of functionally distinct solutions for each condition is shown in Figure 6.4, and the average novelty of the designs for each condition is shown in Figure 6.5. Error bars in Figures 6.4 and 6.5 represent the standard error. Participants in all conditions answered an average of 88% of the post-experiment quiz questions correctly and this percentage did not differ significantly between conditions; thus any observed differences were not due to a failure to encode and access the presented information. For all statistical tests an α level of .05 was used.

Table 6.1 – Average number of ideas per participant in each category

	Pendulum period	Swinging rate of decay	Rate of drip (counting drip sounds)	Rate of flow/fill (amount of liquid)	Rate of burn	Unwinding or pulling of tape	Rate of heating or cooling	Rate of paint drying
Control	0.44	0.11	0.26	1.20	1.50	0.06	0.67	0.11
Clocks-before	0.74	0.06	0.71	0.71	1.62	0.24	0.74	0.12
Devices-before	0.94	0.00	0.22	1.61	2.03	0.17	0.89	0.28
Devices-during	0.61	0.06	0.25	1.67	1.44	0.33	0.89	0.33
Average	0.68	0.06	0.36	1.30	1.65	0.20	0.79	0.21

	Rate of free fall or inclined rolling (gravity)	Repetitious conscious behavior	Heart rate	Drink the vodka	Weight equilibrium	Rate of evaporation	Rate of spring release	Other
Control	0.28	0.50	0.00	0.17	0.20	0.11	0.00	0.06
Clocks-before	0.56	0.44	0.00	0.12	0.03	0.29	0.18	0.29
Devices-before	0.78	0.06	0.11	0.06	0.00	0.22	0.11	0.19
Devices-during	0.42	0.56	0.06	0.33	0.14	0.36	0.14	0.36
Average	0.51	0.39	0.04	0.17	0.09	0.25	0.11	0.23

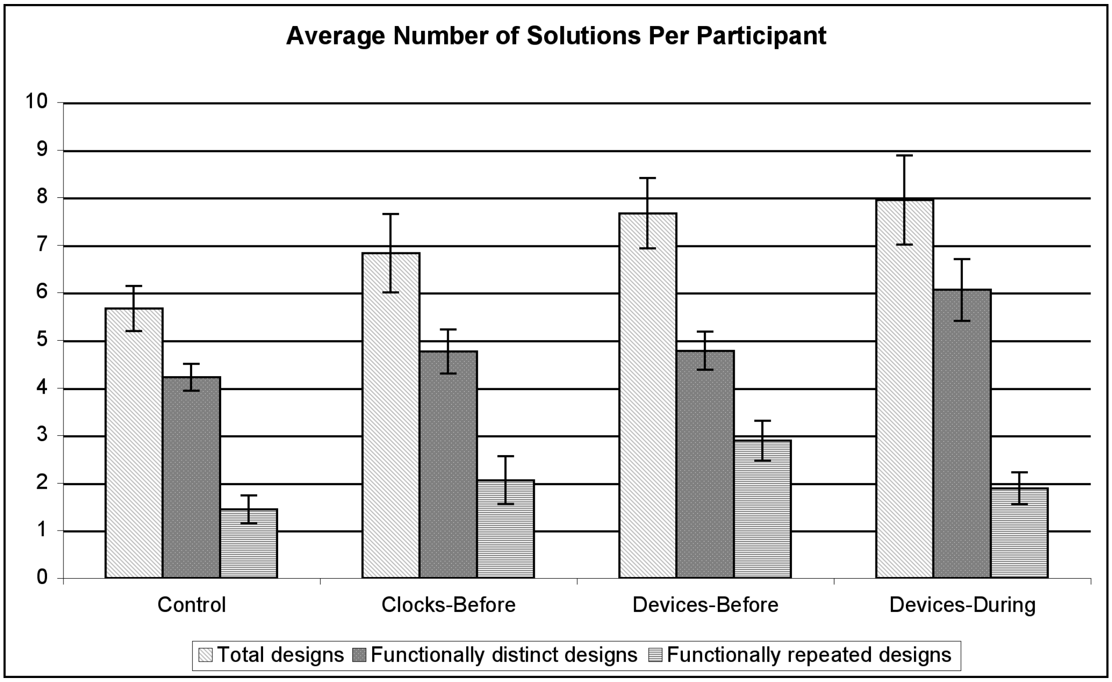

Figure 6.4 – Average number of solutions per participant

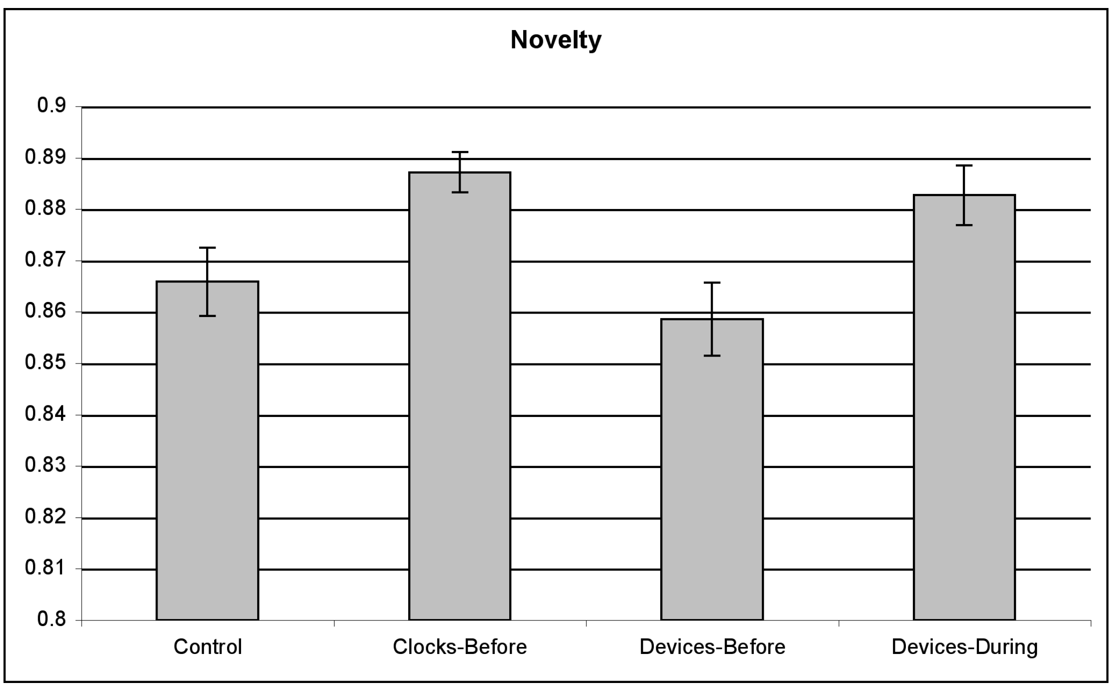

Figure 6.5 – Average novelty per condition

6.1.3.1 Open goals and distantly related devices

The first hypothesis was that devices which are more distantly related to the problem would impact idea generation only when there was an open goal to solve the problem. This hypothesis was examined by comparing the devices-before condition to the devices-during condition.

Participants in the devices-before condition produced more total designs, $t(34) = 2.28$, $p = .03$, than participants in the control condition, but with more functional repeats, $t(34) = 2.92$, $p = .006$. The devices-before condition did not differ significantly from control in the number of functionally distinct designs or novelty. This shows that some of the information was recognized and applied, although with only an increase to the quantity of solutions and not to the variety. Participants in the devices-during condition produced solutions that were marginally more novel, $t(34) = 1.92$, $p = .06$, as well as more total designs, $t(34) = 2.11$, $p = .04$, without the added functional repeats, resulting in more functionally distinct designs, $t(34) = 2.50$, $p = .02$, when compared to the participants in the control condition.

Participants in the during-devices condition generated fewer functionally repeated solutions, $t(34) = 2.03$, $p = .05$, solutions that scored higher in novelty, $t(34) = 2.63$, $p = .01$, and marginally more functionally distinct solutions, $t(34) = 1.87$, $p = .07$, than participants in the devices-before condition. To investigate the timing issue in more detail, the number of functionally distinct solutions in the pre-break period and the post-break period were examined for these two conditions. Participants in both conditions produced a similar number of functionally distinct solutions in pre-break problem solving, $t(34) = 0.14$, $p = .89$, but in post-break problem solving, participants in the devices-during condition produced significantly more functionally distinct solutions than participants in the devices-before condition, $t(34) = 2.54$, $p = .02$. In other words, the distantly related set of three distant device descriptions presented before the problem did not give the devices-before condition any advantage in the pre-break period, but the distant device descriptions did give the devices-during condition a significant advantage in the post-break period.

While there was some effect on the devices-before condition relative to control, comparing the devices-before and devices-during conditions clearly shows that the device descriptions affected problem solving significantly more when there was an open problem solving goal. These results support the first hypothesis that having an open goal increases the positive effect of distantly related information.

6.1.3.2 Surface similarity

The second hypothesis was that information which is more closely related to the problem would impact idea generation more than distantly related information when both were given before the problem began. This hypothesis was evaluated by comparing the clocks-before and devices-before conditions.

These two conditions did not differ significantly in terms of the total number of solutions generated, the number of functionally repeated solutions, or the number of functionally distinct solutions. The level of surface similarity of the material did affect the novelty of the solutions generated as participants in the clocks-before condition scored significantly higher in novelty than participants in the devices-before condition, $t(34) = 3.46$, $p = .002$. Participants in the clocks-before condition also scored significantly higher in novelty than participants in the control condition, $t(33) = 2.73$, $p = .01$. As stated earlier, the participants in the devices-before condition generated more solutions in total, but failed to generate more functionally distinct solutions or solutions high in novelty than the control condition. These results support the second hypothesis that information that is more obviously similar to the problem impacts idea generation more than distantly related information when seen before problem solving has begun. However, this highly related information only impacted the novelty of the solutions and none of the other measures. The analyses in the next section shed some light onto why the novelty of the solutions increased in the clocks condition relative to the control and devices-before conditions.

6.1.3.3 Priming of design solutions

The third hypothesis was that the information provided to participants was expected to prime specific functional principles to be used in solving the problem. The three clock descriptions were expected to prime pendulum based solutions, and the three distant device descriptions, a heart rate monitor, a water flow meter, and a cassette tape deck, were expected to prime solutions in the heart rate, rate of flow/fill, and unwinding and pulling of tape categories respectively. As can be seen in Table 6.1, all four primed functional principles did appear more frequently in the solutions generated by participants in corresponding conditions than solutions generated by participants in the control condition.

Since many of the participants did not produce designs in a particular category, a non-parametric test, the Wilcoxon rank sum test, was used to assess the priming effects rather than a t-test. Participants in the devices-during condition generated marginally more rate of unwinding and pulling of tape solutions than the participants in the control condition, $W=125.5$, $p = 0.078$. The three obtained heart-rate solutions only occurred in conditions where participants were presented with the distant device descriptions. It is difficult to measure the statistical significance of this result due to the low frequency with which it occurred. The devices conditions generated more rate of flow/fill solutions than the control condition, and the clocks condition generated more pendulum solutions than the control condition, although these expected priming effects did not approach or reach statistical significance.

An unexpected finding that was noticed while examining the distribution of solutions across categories was that there was an inverse relationship between two functional categories. The clocks-before condition produced fewer rate of flow/fill solutions than were produced by the participants in the other conditions while producing more drip solutions than the other conditions. This is interesting because both types of solutions involve measuring a quantity of liquid as it leaves a container. In the flow/fill solutions the amount of liquid flowing into or out of a container is used to measure time while in the drip solutions the number of drips as the liquid flows is counted.

To test whether this tradeoff between the two categories was significant, a preference score was created for each participant in the clocks-before and control conditions by subtracting the number of drip solutions from the number of flow/fill solutions. Participants in the clocks condition had a higher preference score for drip solutions more than flow/fill solutions, $W = 91.5$, $p = 0.04$. Individual analysis of the participants shows that generating a solution in either the rate-of-drip or rate-of-flow/fill categories seems to prevent the participant from generating any solutions in the other category. Out of 71 participants only 12 generated both rate-of-flow/fill solutions and rate-of-drip solutions. Of those 12, eight of them switched from one to the other at the break, and only four switched during a problem solving session. The participants that first generated a rate-of-flow/fill solution later generated a total of 82 rate-of-flow/fill solutions and only 4 rate-of-drip solutions. Similarly of the participants that first generated a rate-of-drip solution would go on to generate a total of 22 rate-of-drip solutions and 12 rate-of-flow/fill solutions. This is evidence that the participants were fixated on one problem solution category, which then prevented them from generating solutions in the other category. One possible explanation for this fixation was that the tick-tock noise described in one of the clock descriptions primed the rate-of-drip solutions. Another possibility is that the clocks descriptions primed measuring a liquid in a discrete counting drips way which then inhibited thinking about measuring the liquid in the continuous flow/fill way.

6.1.4 DISCUSSION

The results support all three hypotheses. There was strong support for the hypothesis that open problem solving goals influence the acquisition and use of distantly related information. The results also agree with prior work on analogical transfer showing that distantly related information is often not recognized as relevant, but that information that shares surface similarity with the problem is recognized as relevant. There was also evidence that the functional principles in the presented devices were primed and used in the solutions.

Open problem solving goals have been shown to influence information acquisition in problem solving even when people are not working on a problem (Moss et al., 2007). However, this initial work on open goals used simple problems. The results presented here extend this work to a more complex problem and indicate how open goals may interact with analogical transfer by allowing for the recognition and use of distantly related analogies.

When devices that were functionally related to the problem but not related in purpose or appearance were presented before participants had a chance to attempt the problem, it was indeed difficult for participants to recognize and apply the information, resulting in no more functionally distinct solutions than from participants who received no relevant information. This same distantly related information, presented after the participants were given five minutes to work on the problem, resulted in a significant increase in both the number of functionally distinct solutions and the novelty of the solutions.

Participants who received the priming examples generated more solutions in all primed solution categories than participants in the control condition. Although this effect did not always reach statistical significance, all four primed examples saw shifts in the number of solutions in the correct direction when compared to the control case. The incorporation of aspects of example solutions has been shown to occur in design (e.g., Jansson & Smith, 1991; Purcell & Gero, 1996), and so it is not surprising that we found them as well. What is interesting is the extent to which distantly related devices primed solution concepts. Most prior work on design idea generation has focused on presenting examples that are actually solutions to the problem at hand (e.g., Jansson & Smith, 1991; Perttula & Liikkanen, 2006). Our results therefore extend this prior work by showing that distantly related information can actually prime solution concepts when presented during a break in problem solving. The optimal timing of such information is left for future work.

The clock descriptions bear more surface similarity to the problem since they are literally time keeping devices, so analogies from them can be more easily applied to problem solving. Because of this, participants who received the clock descriptions before starting the problem scored significantly higher in novelty when compared to the participants who received device descriptions before starting the problem. The clocks conditions apparently primed the creation of drip counting solutions which were less frequent in the other three conditions, and therefore these solutions increased the novelty of the clocks-before condition because the solutions generated by participants in the clocks-before condition were appreciably different from the solutions generated by participants in the other conditions. This priming could have occurred because the clocks primed thinking about measuring time discretely as counting the number of drips or because the tick-tock in the clocks primed the sound of dripping. This change in the distribution of solutions and the lack of differences between control and the devices-before condition is evidence that in the absence of open goals, surface similar information is more readily applied to problem solving than surface dissimilar information.

The inhibition of one or more solutions caused by a block or fixation on prior ideas is a common theme in the problem solving literature (e.g., Duncker 1935/1945; Janson & Smith, 1991; Smith & Blankenship, 1991). When a problem solver starts a problem, it may initially be easy to generate different ideas, but after generating a few ideas it becomes harder to generate new ideas because the previously generated ideas interfere with the ability to generate future ideas. This kind of fixation has been shown in simple problems (Moss et al., 2008). In computational models of human memory such as ACT-R (Anderson et al., 2004) an item's probability of retrieval is based on how well a person's current context primes the item as well as how recently and frequently the item has been retrieved in the past. So, the first few ideas that a person generates in a design problem may be continuously retrieved both because these were the ideas that were best primed by the given problem and because they have recently been retrieved. One approach to overcoming this fixation is to take a break from the problem. This helps to overcome fixation due to frequency and recency of retrieval, but it does not change the long-term associations that led to the initial solution concepts in the first place (Wiley, 1998). However, exposure to new information after there is an open problem solving goal may allow new ideas to enter the problem solving process and help to overcome fixation due to long-term associations as a proposed model of the open goal effect states that open goals lead to the formation or strengthening of associations between the problem and relevant information that is encountered after a problem solving goal has been established (Moss, 2007).

These results have a number of implications for improving design methodology. Analogical inspiration in design can clearly be a powerful way to increase the number and variety of solutions generated in problem solving leading to better and more novel designs. From the results of this experiment and from prior research, the best time to seek analogical inspiration for maximum effect is after work on the problem has begun. In fact, the point at which the designer reaches an impasse in problem solving, namely when no new significant design concepts are being generated, may be the best time to take a break (Moss et al., 2008). When searching for analogical inspiration, both information that is surface similar and dissimilar to the problem solving task at hand can be considered, resulting in the possibility of wide variation in potentially inspirational information, but the dissimilar information is the most influential and effective when received after problem solving has begun. Given that it has been found that it is not necessary for the problem solver to even be aware of encountering the relevant information for it to have an impact on problem solving, it may be best to engage in a variety of tasks where exposure to disparate information is encountered. There are even opportunities for design tools that aid idea generation by presenting a wide variety of design stimuli since people are generally not very good at coming up with distant analogies on their own.

6.1.5 CONCLUSIONS

The timing and analogical similarity of newly acquired information plays a role in generating ideas and solving problems in design. By manipulating the type and timing of relevant information, it was found that highly similar information impacted problem solving even before problem solving began, but distantly related information only affected problem solving when it was presented during a break. These results support the idea that open goals increase the likelihood that distantly related information become incorporated into problem solving. These distantly related ideas might help to inspire innovative or creative solutions to design problems.

Functional principles found in the problem-relevant information given were found to prime solutions in corresponding categories. Evidence of induced design fixation was observed as participants exhibited an interesting tradeoff behavior when thinking about two distinct solution approaches (liquid flowing versus counting the number of drips). This relationship suggests that the participants became fixated from a priming hint, and were unable to generate solutions from the other similar solution category.

In solving design problems that incorporate multiple objectives, such as stylistic form and function, designers need to be able to apply new knowledge about the problem or about relationships between different objectives in the problem, which can be highly beneficial to the effective management of tradeoffs or exploitation of synergies between the multiple objectives. This experiment helps to demonstrate the power of using analogical inspiration to help designers apply knowledge that can benefit problem solving, as well as suggestions of when and how to present the information for maximum benefit. In the next section an experiment is conducted to use the findings of this experiment to present analogical inspiration to help participants solve a design problem that involves both functional and form goals.

6.2 INSPIRING DESIGNERS OF FORM AND FUNCTION DESIGN PROBLEMS

Searching a multi-objective design space for an optimal solution can be a daunting task. Any information that can benefit problem solving would certainly be appreciated, but as seen in the previous experiment, having access to beneficial information is not always enough to allow the designer to assimilate it into the problem-solving process. More specifically, it was shown that information that bears little surface similarity to the problem being solved could be difficult to incorporate, and that if incorporated, a wider range of relevant materials is most beneficial. It was also shown that actively working on the problem and being familiar with the goals of the problem can help increase a designer's sensitivity to beneficial information, increasing the likelihood that the information will benefit problem solving. Much research has shown that designers are particularly sensitive to example solutions (Jansson & Smith, 1991; Purcell & Gero, 1996; Chrysikou & Weisberg, 2005; Perttula & Liikkanen, 2006), but that caution must be exercised since the effects of example solutions can cause fixation on negative elements in the example design (Jansson & Smith, 1991; Linsey et al., 2010). This problem is further compounded by the fact that designers, even experienced designers, tend not to be aware that they are fixated (Ward, 1994; Kolodner, 1997; Marsh et al., 1999; Linsey et al., 2010). Conversely it has also been shown that presenting designers with a variety of different example solutions and alternate representations (Linsey et al., 2010) can negate the effects of fixation. Presenting designers with good example solutions have also been shown to benefit problem solving in design teams (Fu et al., 2010), and visual analogies have been shown to be beneficial for both novice and expert architects in problem solving (Casakin & Goldschmidt, 1999).

The study presented in this section extends the findings of the experiment described in Section 6.1 to a form and function design problem in order to examine the use of visual examples as a tool to help designers better incorporate beneficial information in design problem solving. More specifically, participants were asked to discuss relationships about tradeoffs and trends in vehicle design that were expected to be beneficial and inspirational in solving the multi-objective vehicle design problem. This task was given immediately after the participant received and reviewed the problem statement and was familiar with the process of solving the problem and had opened problem-solving goals. While discussing the problem-relevant and inspirational information, participants had access to information for one of three conditions: 1) images of three dissimilar real cars, 2) three similar real cars, or 3) no images at all. In conditions where participants had access to car images, the participants were encouraged to identify elements of the vehicle designs to demonstrate the relationships that benefit or hinder vehicle performance.

6.2.1 HYPOTHESIS

Two hypotheses were examined in this experiment: 1) Problem-relevant information presented in the form of a discussion would benefit problem-solving more if presented with analogous visual examples, and 2) a wider breadth of example designs would be more beneficial to problem solving than a narrower breadth of example designs.

The first hypothesis assumes that these visual examples will help to demonstrate the applicability of theoretical relationships in vehicle design with the problem-solving task itself. By demonstrating this applicability, participants should be better equipped to apply this information effectively and efficiently to problem solving. The second hypothesis pulls from the experiment in Section 6.1 and from the literature (Linsey et al., 2010) in assuming that a wider range of example solutions will give the participants easier access to the information. More specifically, participants will be given the opportunity not only to discuss why a certain design was chosen, but the variety in car designs depicted gives them an opportunity to compare and contrast different vehicle shapes. This benefit does not apply when comparing three similarly shaped car designs.

6.2.2 METHODS

Two car design problems were administered serially in this experiment. All participants, regardless of condition, solved the same two design problems. Both car design problems were adapted from the illustrative examples presented in Section 5.4.2, which in turn were inspired by the tradeoffs identified in Section 4.1. In the first car design problem, participants were asked to design a sports car that can be used every day, which is one that rates above 90% in sportiness, has an aerodynamic coefficient of drag below 0.25, and interior volume of over 75 cu ft. This first design problem is shown in Figure 6.7. In the second car design problem, participants were asked to design a car that combines excellent fuel efficiency, a high driving position, and large interior volume. These goals were then translated to target values of an aerodynamic coefficient of drag below 0.25, a center of gravity above 25 inches, and an interior volume above 83 cu ft. This second design problem is shown in Figure 6.8.

The purpose of the first car design problem was twofold. First, because all participants, regardless of condition, were given the same information to solve the same problem, the measured performance of this first problem helped to ensure that the relative performance of each condition was statistically the same. The second reason was to provide participants a chance to practice and familiarize themselves with the interface and to open goals pertaining to the relationships between aerodynamic performance and interior volume.

The purpose of the second car design problem was to measure the effects of the inspirational material received with the second design problem, which was dependant on which condition the participant was assigned to. Three conditions were tested in this experiment, a No Example condition, a Similar Examples condition, and a Dissimilar Examples condition. In the No Example condition, participants were given the blurb shown in Figure 6.9 immediately after the problem statement for problem 2, as shown in Figure 6.8. No vehicle design images are included for participants of the No Example condition. In the Similar Examples condition, participants received the blurb and three real world images of similar cars as shown in Figure 6.10. In the Dissimilar Examples condition, participants received the blurb and three real world images of dissimilar cars as shown in Figure 6.11. These three conditions are analogous to the similarity and dissimilarity of the clocks and devices conditions in the experiment described in Section 6.1.

The concept of similarity and dissimilarity of vehicle designs in this experiment pertains to the silhouette of the vehicles and what parametric values would be used to model the three designs. Despite belonging to three different classes of vehicles, all three similar vehicle designs feature curvaceous designs with raked window designs, smooth roofline profiles, raked noses, short trunk profiles, and a sloping rear end. The first car image shown in the Dissimilar Examples condition is the same as the one shown in the Similar Examples condition, a 2009 Toyota Prius. The two other vehicles shown are dissimilar from the Prius. The second vehicle is a 2011 Range Rover Sport, which has a boxier design with windscreens and a nose that are more vertical, a higher ground clearance, sharp roof transitions, and a less sloping rear end. The third vehicle is a 2011 Dodge Challenger, which features a very long hood, a more vertical nose, sharp roof transitions, and a long trunk profile. In order to assess these differences more rigorously, two independent people who did not participate in the experiment were asked to use the design interface in this experiment to design cars that best resembled each of the five different vehicles that were presented as visual inspiration. The three designs created to resemble the vehicles presented in the Similar Examples condition are shown in Figure 6.6(a), (c), and (d), and the three designs created to resemble the vehicles presented in the Dissimilar Examples condition are shown in Figure 6.6(b), (d), and (f). A comparison between the average squared differences between each of the parameter values of each pairing of cars in each condition is shown in Table 6.2, where a higher number indicates a larger difference between the parameters used to draw the cars, and thus a higher level of dissimilarity. As can be seen, the differences between the three cars presented in the Similar Examples condition are substantially smaller than the differences between the three cars presented in the Dissimilar Examples condition.

Figure 6.6(a) – Similar Designs #1 (Toyota)

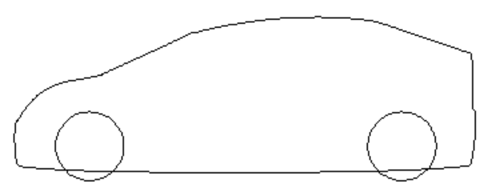

Figure 6.6(b) – Dissimilar Designs #1 (Toyota)

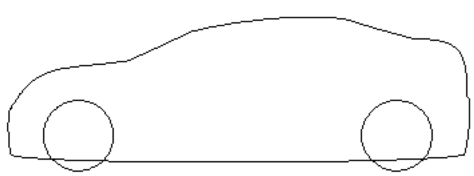

Figure 6.6(c) – Similar Designs #2 (Hyundai)

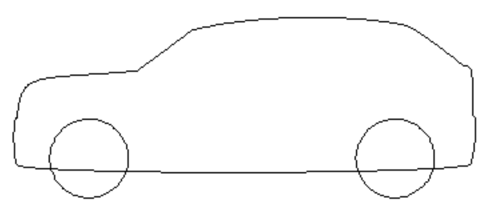

Figure 6.6(d) – Dissimilar Designs #2 (R.Rover)

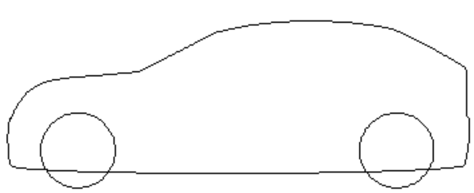

Figure 6.6(e) – Similar Designs #3 (Infiniti)

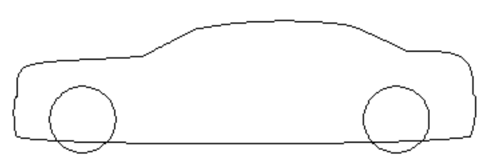

Figure 6.6(f) – Dissimilar Designs #3 (Dodge)

Table 6.2 – Average Squared Differences vehicles in Similar Examples and Dissimilar Examples conditions

Average Squared Difference between Similar Examples		Average Squared Difference between Dissimilar Examples	
Design 1 & Design 2	656.9	Design 1 & Design 2	2352.6
Design 2 & Design 3	956.5	Design 2 & Design 3	2648.4
Design 1 & Design 3	997.8	Design 1 & Design 3	2300.2
Average Difference	**870.4**	**Average Difference**	**2433.7**

To measure the expected effects, the problem solving performance of the three conditions were measured by recording the number of iterations needed by each participant to complete the second design problem. A lower number of iterations needed to find the solution suggests that the participant was better able to utilize the information given to efficiently solve the problem.

6.2.2.1 Participants

The same eighteen participants from the Artificial Neural Network Learning and Generation experiment in Section 5.3 also participated in this experiment. The eighteen participants were randomly separated into three equal conditions of six participants each. Because all participants had participated in the same experiment prior to this experiment, and because that experiment did not involve separating the participants into different conditions, their additional experience would not favor any condition more than any other.

6.2.2.2 Design Interface

A vehicle design interface was developed in MATLAB for this experiment. The starting point of the design interface for Problem 1 is shown in Figure 6.12, and the starting point of the design interface for Problem 2 is shown in Figure 6.13. The interface displays a vehicle design that is generated using the parametric model detailed in Section 3.1 The interface features twelve labeled sliders that control the twelve parameters in the parametric model. Performance statistics for the current vehicle design are provided in the lower right of the interface for each of the three goals for each problem. The interface for problem 1 displays the Aerodynamic Performance of the vehicle in cD, the Volume in cubic feet, and the Sportiness Rating in percent. The interface for problem 2 displays the Aerodynamic Performance in cD, the Center of Gravity Height in inches, and the Volume in cubic feet.

The sportiness rating used in the first car design problem uses an Artificial Neural Network developed using the methods described in Section 5.3. This Artificial Neural Network was trained using data from an independent rater who did not participate in this study. The aerodynamic coefficient of drag estimates were provided using the parametric model detailed in Section 3.2, and the interior volume and center of gravity estimates were provided using the program developed in Section 3.3.

The interface starts with all sliders in the middle position as shown in Figures 6.12 and 6.13. After the participant selects the desired parametric values for the next iteration of their car design, they click the button labeled Generate Car in the lower right corner. This button updates the image of the car based on the value of the sliders, and updates the values of the three performance statistics. The interface program outputs a text file that reports the parameter values and performance statistics for all design iterations generated by each participant.

6.2.2.3 Procedure

Before starting, each participant received a quick primer for how to use the interface and what each vehicle design parameter affects. At the start of each problem solving session, a paper copy of the respective problem statement was given to the participant and read aloud to them. Each participant was asked if they had any questions before they began each problem solving session. Participants were encouraged to design the vehicles by looking at the vehicle images, and not to focus exclusively on the slider positions. Participants were assured by the researcher that a solution does in fact exist for both problems. Participants were allowed to refer to the problem statement handout and inspirational images freely during problem solving.

Each participant was given 25 minutes to solve the first design problem, and unlimited time to solve the second design problem. The first problem took an average of 21 minutes for participants to solve, with the shortest taking 9 minutes. Five participants were unable to finish the first design problem within the allotted time. The second problem took participants an average of 16 minutes to solve, with the shortest taking 6 minutes, and the longest taking 38 minutes.

After completion of the first design problem, the participant was given a chance to take a break. After the break, the second problem statement was given, and the inspirational material was administered. After reading the problem statement to the participant, the blurb that corresponds to their condition was also read to them. In relevant conditions, the three vehicle design images were then shown to the participants. All participants were then asked to identify for the researcher all relevant vehicle design characteristics that would help to achieve the goals stated in the problem statement for problem 2. In the Similar Examples condition, participants were then asked to identify which design elements were designed into the vehicles presented to achieve the problem goals and to discuss the overall designs of those vehicles. In the Dissimilar Examples condition, participants were asked identify which design elements were designed into the vehicles presented to achieve the problem goals and to compare and contrast between the different goals of each vehicle, and how various characteristics contribute to those goals, and how they differ from each other.

In all conditions, participants were corrected if the trait or feature they discussed as beneficial for a design goal was in fact not beneficial, or vice versa. All participants were asked to continue listing characteristics until the participant mentioned low vehicle height or profile, smooth shapes or smooth transitions, raked front and rear windscreens, and a raked nose. In some cases if more than approximately 3 minutes had elapsed in the discussion without all of the aforementioned characteristics being mentioned, the researcher would ask targeted questions such as "What about the nose?" or "Can you think of any other characteristics about the roof that might benefit aerodynamics performance?" in an attempt to elicit the desired responses. The goal of these intervening questions was to ensure that all participants, regardless of condition, were made aware of and fully understood the design characteristics before moving on to problem solving.

Problem 1 (25 Minutes)

Your goal is to design a sports car that can be used everyday. The threshold for a good sporty looking vehicle is one that ranks higher than 90% in sportiness. In order to be deemed useful, good fuel efficiency, and thus aerodynamics and a capacious interior are also a must. As a result the aerodynamic coefficient of drag must be below 0.25 and the interior volume must exceed 75 cu ft.

Aerodynamic cD lower than 0.25

Volume bigger than 75 cu ft

Sportiness rating above 90%

Figure 6.7 – Statement of design problem for Problem 1

Problem 2

Buyers want a vehicle with excellent fuel efficiency, a high driving position, and large interior volume. Your goal is to design such a vehicle. In order to meet these goals, your design must have an aerodynamic coefficient of drag below 0.25, a center of gravity above 25 inches, and an interior volume above 83 cu ft.

Aerodynamic cD lower than 0.25

Center of gravity taller than 25 inches

Volume bigger than 83 cu ft

Figure 6.8 – Statement of design problem for Problem 2

Identify and discuss the vehicle design traits that might help your vehicle achieve these goals.

Figure 6.9 – Blurb as used in the No Examples condition

Three examples of vehicles from different segments of the market that have been designed to satisfy similar goals have been included on the following page. Identify and discuss which design elements in each of these car designs might have been chosen by their designers to meet each of the goals of your design problem, and which design elements might be detrimental to the performance of each design.

2009 Toyota Prius

2011 Hyundai Sonata

2010 Infiniti FX45

Figure 6.10 – Blurb, and three similar examples as used in the Similar Examples condition

Three examples of vehicles from different segments of the market that have been designed to satisfy similar goals have been included on the following page. Compare and contrast which design elements in each of these car designs might have been chosen by their designers to meet each of the goals of your design problem, and which design elements might be detrimental to the performance of each design.

2009 Toyota Prius

2011 Range Rover Sport

2011 Dodge Challenger

Figure 6.11 – Blurb, and three dissimilar examples as used in the Dissimilar Examples condition

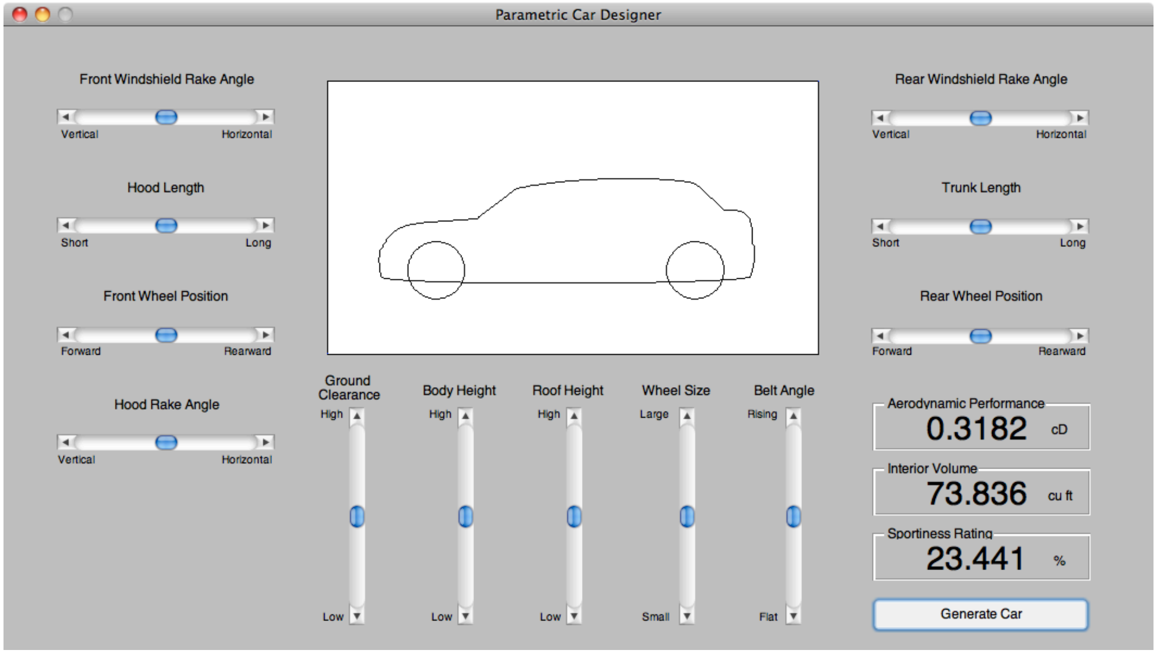

Figure 6.12 – MATLAB vehicle design interface with initial design for Problem 1

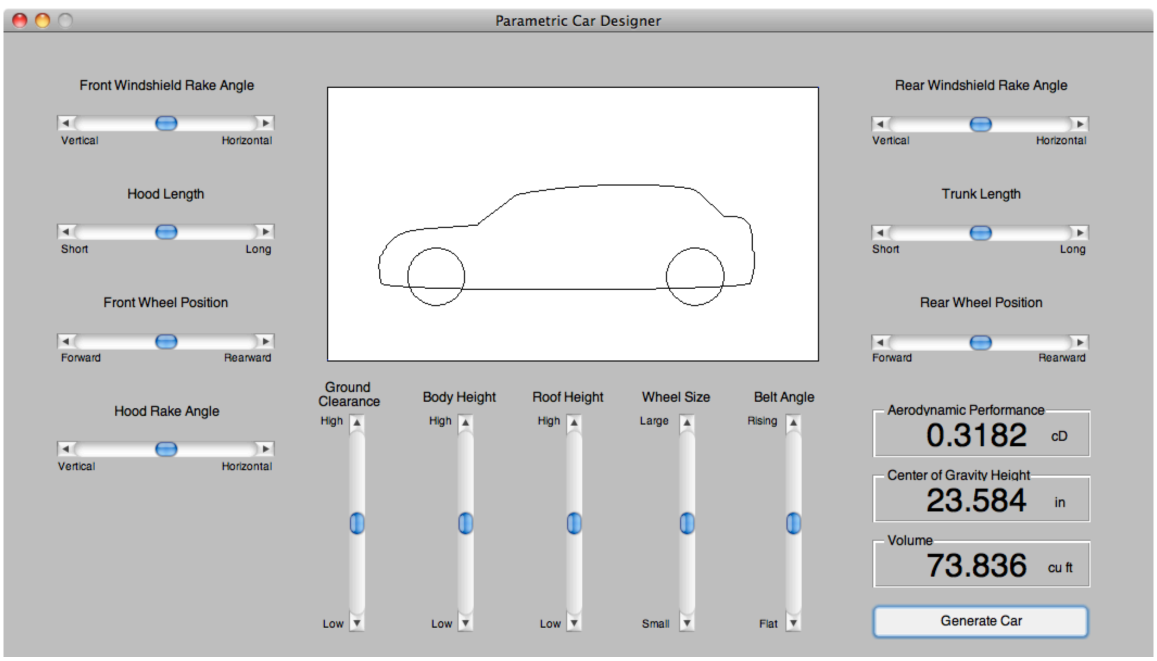

Figure 6.13 – MATLAB vehicle design interface with initial design for Problem 2

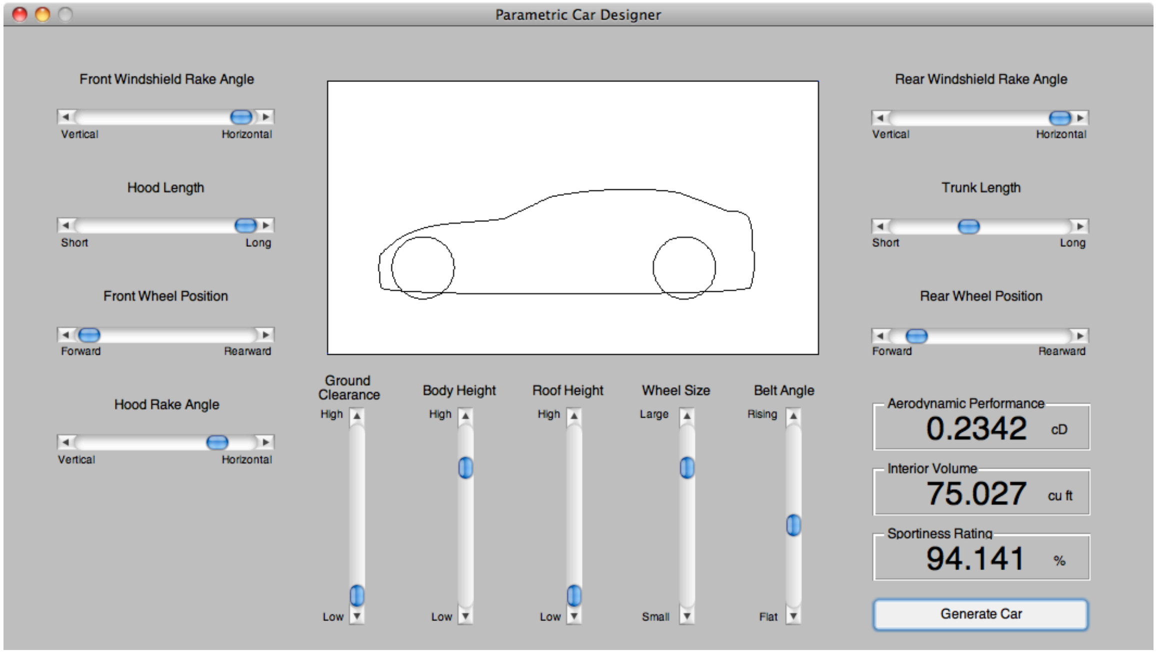

Figure 6.14 – MATLAB vehicle design interface with a valid solution for Problem 1

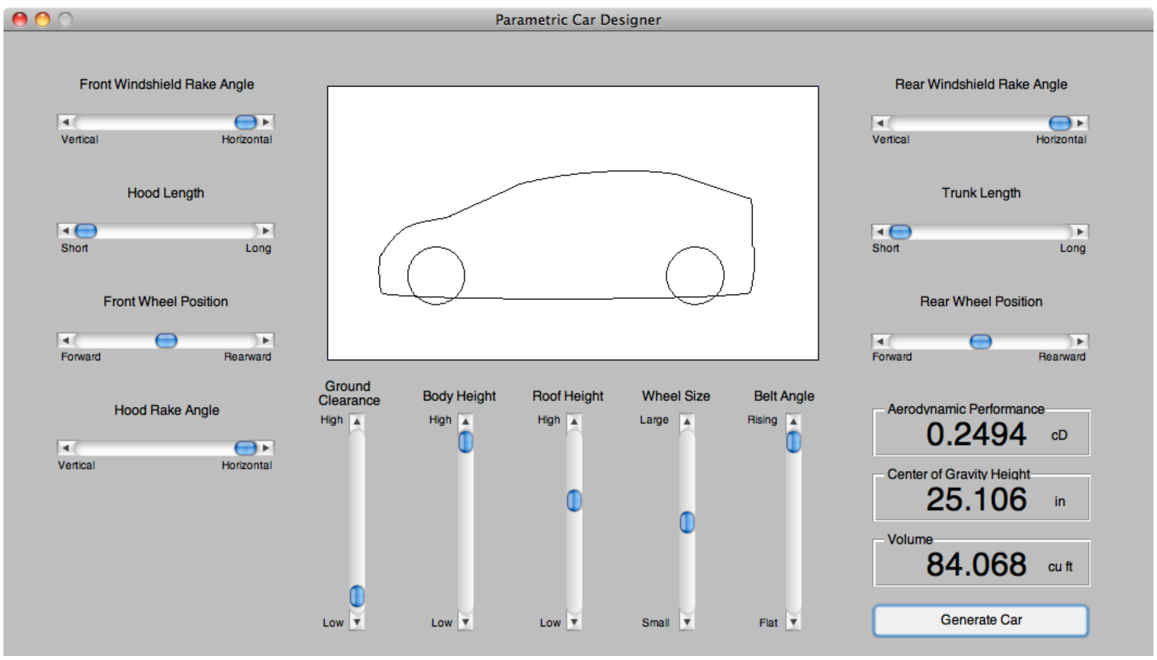

Figure 6.15 – MATLAB vehicle design interface with a valid solution for Problem 1

6.2.3 RESULTS

A series of t-tests were done to compare the average number of iterations per participant between conditions. For all statistical tests, an alpha level of 0.05 was used (α=0.05) to determine significance, and an alpha level of 0.1 was used (α=0.1) to determine marginal significance. The average number of iterations per participant to solve each problem for each condition is shown in Figure 6.16. The error bars in Figure 6.16 represent the standard error.

The first design problem was identical across all participants and serves as a test to ensure that participants were suitably similar in performance across all conditions before the inspirational information was given. As would be expected, there was no statistical difference between the average number of iterations per participant in any pairing of conditions. Comparing the No Examples condition to the Similar Examples condition, $t(10)=0.82$, $p=0.215$, comparing the No Examples condition to the Dissimilar Examples condition, $t(10)=0.536$, $p=0.302$, and comparing the Similar Examples condition to the Dissimilar Examples condition, $t(10)=1.21$, $p=0.126$.

The first hypothesis was that problem-relevant information presented in the form of a discussion would benefit problem-solving more if presented with analogous graphical examples. This hypothesis was examined by comparing the number of iterations per participant for problem 2 between the No Examples condition and both of the Examples conditions. Participants in the Similar Examples condition required significantly fewer iterations to solve the second design problem when compared to participants in the No Examples condition, $t(10)=2.23$, $p=0.025$. Similarly, participants in the Dissimilar Examples condition also required significantly fewer iterations to solve the second design problem when compared to participants in the No Examples condition, $t(10)=3.63$, $p=0.002$.

The second hypothesis was that a wider breadth of example designs would be more beneficial to problem solving than a narrower breadth of example designs. This hypothesis was examined by comparing the number of iterations per participant for problem 2 between the Similar Examples condition and the Dissimilar Examples condition. Participants in the Dissimilar Examples condition required marginally fewer iterations to solve the second design problem than participants in the Similar Examples condition, $t(10)=1.40$, $p=0.096$.

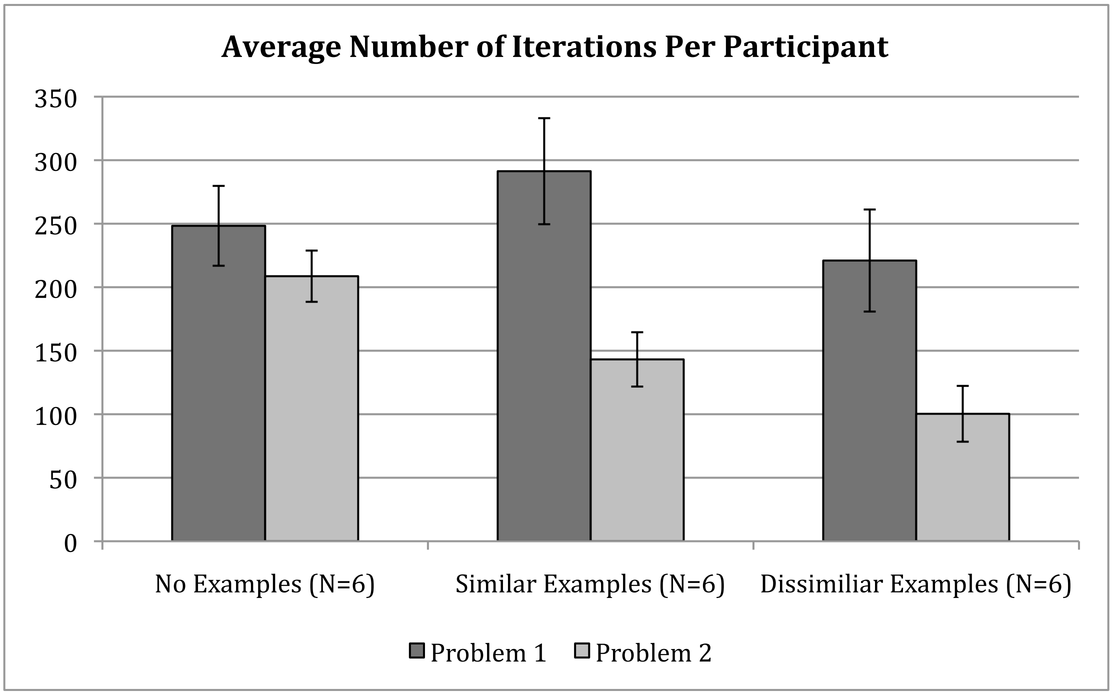

Figure 6.16 – Average number of iterations per participant

6.2.4 DISCUSSION

The results from the experiment in Section 6.1 found that surface similar information is easier for participants to apply to design problems than surface dissimilar information, and that open problem solving goals should help participants apply problem-relevant information more easily to problem solving. This experiment presented participants with surface similar information after problem solving goals were opened. The information presented was in the form of a discussion, where all participants, regardless of condition, were asked to discuss information that is beneficial for solving the problem that came after. Despite the fact that all participants possessed this beneficial information going into the problem solving stage, participants in conditions where vehicle images were given significantly outperformed participants who did not receive the vehicle images. Additionally it was found that presenting participants with a wider range of vehicle designs resulted in marginally better performance than participants who were presented with a narrower range of vehicle designs. It is possible that the use of these images in the discussion provided some additional connection between the information being discussed and the ability to apply the information to the design problem, and that a wider range of vehicle designs increased the applicability of the inspiration by providing even more connections to the design problem. Since there was no condition where participants did not discuss the material before problem solving, it is unknown whether the performance of the No Examples condition is indicative of not having the information at all, or if their performance was improved by the information, but not by as large of a margin as when vehicle images were presented with the information.

The literature on the priming of design solutions (e.g., Jansson & Smith, 1991; Purcell & Gero, 1996, Perttula & Liikkanen, 2006) may provide an alternative explanation for the improved performance for participants who were shown the vehicle images. One of the vehicles shown in both conditions is a 2006 Toyota Prius, which is a vehicle that shares many of the goals of the design problem and also shares many design characteristics with the final solution of the design problem. It is possible that seeing this vehicle prior to problem solving simply primed a Prius-like solution.

6.2.5 CONCLUSIONS

In short, the results of this study strongly support the first hypothesis in that problem-relevant information presented in the form of a discussion was shown to benefit problem solving more when presented with analogous visual examples. The second hypothesis is also supported in that a wider breadth of example designs provided was marginally more beneficial to problem solving than a narrower breadth of example designs.

Additional research is needed that examines in greater detail how and why the vehicle images allowed participants to more efficiently and effectively apply the information to problem solving. This future research should look into the effects of solution priming, and whether the act of applying the knowledge to vehicle images during the discussion provided a deeper connection to the participant's mental representation of the design problem.

CHAPTER SEVEN
CONCLUSIONS, CONTRIBUTIONS, AND FUTURE WORK

7.1 DISSERTATION SUMMARY

This dissertation has leveraged several powerful technologies in the pursuit of unifying stylistic form and function. The long-term purpose is to unify these rival goals into a rich and informative relationship where their differences can be ameliorated and their synergies exploited. This dissertation identified two main limitations that stand in the way of this relationship and has made inroads towards bridging them.

The first of these limitations is a representational barrier. Those managing functional goals have long enjoyed the benefits of computational modeling, numerical optimization, and easily quantifiable goals. Because stylistic form is not as easily quantified, these same benefits are not afforded to those who manage stylistic goals. As a partial result of this representational barrier, the design of stylistic form and the design of functional goals are frequently performed by two independent groups of people using vastly different tools and methods. This dichotomy between stylistic form and function has led to a lack of interface between these goals, which can often result in an underutilization of potential synergies and poor management of tradeoffs.

In an attempt to reduce the representational barrier between goals of stylistic form and function, the research in this dissertation has adopted technologies from computer science, mathematics, machine learning, psychology, decision sciences, and statistics to explore three methods to model stylistic form. The results of this exploration culminated in the development of an Artificial Neural Network based learning system that allows computers to accurately model human judgments of stylistic form. This model quantifies stylistic form, which allows it to be used in unified frameworks with functional goals. One such framework couples this modeling system with a Genetic Algorithm to enable computers to generate new designs that successfully elicit targeted responses of stylistic form from consumers. This method also breaks down barriers to unify goals of both stylistic form and function in computational design generation. To illustrate form and function design abilities of this Artificial Neural Network and Genetic Algorithm system, a Pareto Set of the trade off between a stylistic form goal and a functional goal was mapped, and several designs that target a variety of stylistic form and function goals were generated computationally.

Before this learning and generating method can be used to interface with stylistic form, a deterministic computer model of the desired product type must first be obtained or built. The work in this dissertation developed a parametric vehicle design model and several functional assessment models for vehicle design to satisfy this need. There are literally a limitless number of product categories that can be modeled and used with the framework discussed here.

The second limitation that prevents the unification of stylistic form and function is the lack of formalized methods that allow human designers to better manage tradeoffs and synergies between stylistic form goals and functional goals in design problems. This problem is exacerbated by research that shows that designers, even engineering design faculty, can becoming fixated on a specific representation of a problem (Jansson & Smith, 1991; Linsey et al., 2010), and that these designers may not be aware of this limited representation of the design problem (Ward, 1994; Kolodner, 1997; Marsh et al., 1999; Linsey et al., 2010).

To help formalize methods that aid human designers to better represent the problem for more efficient management of tradeoffs and synergies, two experiments were conducted towards this goal. The first experiment found that designers might have trouble applying information that they encounter before problem solving if the information does not share surface similarities, such as a shared domain, with the problem being solved. Additionally, it was found that information that is more diverse in nature, while more difficult to apply without open goals, can have bigger benefits to problem solving if the open goals are present to allow the information to be assimilated. The findings of this experiment led to a second experiment that found that while surface similarities and open goals alone can enable designers to better apply inspirational information to a design problem, further improvements can be had by giving designers inspirational visual aids that allow them to graft their knowledge more directly onto the representation of the problem being solved, and that in agreement with the first study, this study showed a marginal effect that a wider range of inspirational information is more beneficial to problem solving than a narrower range.

A common thread that unites the computational and cognitive work in this dissertation is that it all foundationally stems from a better understanding of people as both designers and consumers. In the case of the cognitive studies, a better understanding of how designers apply analogical inspiration to improve problem solving inspired a series of experiments that demonstrate formalizable suggestions that may help to unlock an improved ability to manipulate form and function goals in design problem solving. In the case of the computational modeling system, this unification is made possible by a better understanding of how consumers judge stylistic form and function, which leads to a better understanding of how to survey these consumer judgments, which opens the door for a successful modeling and generation system for stylistic form and functional goals. This modeling and generation system enables stylistic form to be quantified for equivalent treatment with functional goals.

By bridging the limitations in computational representation and cognitive formalization, this dissertation has shown how the hierarchy of stylistic form following function can be unified into a mutual relationship that benefits both designers and consumers.

7.2 CONTRIBUTIONS

The research in this dissertation has made the following contributions to the field of product design.

1. An Artificial Neural Network based framework for modeling human stylistic and functional form judgments which
   a. Quantifies stylistic form and functional form judgments using machine learning models for concurrent optimization with functional goals.
   b. Enables computers to use survey data to predict how consumers will react to whether a design adheres to desired stylistic conventions, or if the design communicates desired functional performance.
   c. Can be coupled with numerical optimization techniques, such as Genetic Algortihms, to computationally generate optimal designs that can simultaneously target multiple consumer desires in stylistic form and functional goals.
   d. Can be coupled with numerical optimization techniques, such as Genetic Algorithms, to computationally generate a Pareto Set of optimal solutions that map optimal tradeoffs between stylistic form and functional goals.
   e. Is able to model stylistic form and functional form judgments in such a way to capture interaction effects between different parameters.
   f. Is able to model continuous parameter variations and provide a continuous learned output.

2. Contribution to the understanding of how people use judgments of functionality and stylistic traits to create judgments on other functional and stylistic traits of a product.

3. Contribution to the understanding of how to inspire designers to more efficiently and effectively solve form and function design problems.

4. Contribution to the understanding of the effect of open goals, timing, and information characteristics for designers to better assimilate new information into design problem solving.

7.3 FUTURE WORK

The research presented in this dissertation provides a good springboard for additional research and future work.

7.3.1 FUTURE COMPUTATIONAL MODELING

The computational modeling and generation framework used in this dissertation can measure and model human judgment and generate designs for a limitless number of product categories. Theoretically, any product category that can be represented with a parametric model that generates a design deterministically based on the values of the parameters can be used with this framework. Future research should study the use of these methods with other product domains to verify that this methodology can truly be extended to these other domains. Once a wider variety of product domains are spanned, and a larger corpus of product categories and survey results are achieved, it may be interesting to determine connections and universal rules that can be learned throughout all product categories. One such example would be to study whether specific universal design characteristics, such as gentle curves, or sharp acute angles have any universal ramifications on human judgments across multiple product categories. Studying these universal ramifications may lead to a deeper understanding of the form language of product shaping, and may be used to develop general shape grammars or form language metrics to define what it may mean for a design in any product category to elicit specific consumer preferences. Researchers (Achiche & Ahmed, 2008) have targeted the idea of determining these universal form language rules by surveying non-product oriented shapes and analyzing these results with fuzzy logic. Some rules were generalized from this study, but it was noted by the researchers that having more applicability to actual products might help people recognize these features more accurately.

The method discussed in this dissertation to model human judgment using Artificial Neural Networks has illustrated the potential to capture what was once nebulous. This research uses a commercially available solution to build, train, and manipulate the Artificial Neural Network. Much research has been done to further improve the accuracy of ANNs (Baxt, 1992; Hashem & Schmeiser, 1995; Zhang, 2000), and improve their performance with sparse or deficient data (Bishop, 1991; Tresp et al., 1994; Twomey & Smith, 2002). Utilizing these refinements with a bespoke ANN platform may help to further improve model accuracy and performance for the needs of capturing and modeling human judgment. In addition to refinements of the ANN itself, future research should examine ways to improve the survey process. Much research has shown impressive capabilities with Convolutional Neural Networks for visual recognition and classification of diagrams and images (Lawrence et al., 1997; Simard et al., 2003; Fu & Kara, 2011). Integrating this visual recognition and classification work with the modeling of stylistic form may negate the tedious and potentially limiting requirement of using a representative computational model. Instead, it may be possible to use images of products, whether fictitious or real, directly to gauge the stylistic form desires of consumers. These improvements could help to improve both the capabilities of measuring human stylistic desires, and the process used to model them, although the process needed to generate new designs from what is learned from this method is then unclear. Other possibilities that may improve the survey techniques employed by this framework involve further research into improving the type of questions and stimuli used to gauge participant responses. Future research can use Principal Component Analysis, as demonstrated in Section 4.1.2.5 with data from a large number of potential stylistic form descriptors in order to determine a more optimal set of descriptors that are highly independent from each other to more efficiently and effectively gauge consumer preferences in future work. Also, more research should be done to examine the optimal number of Likert Scale levels to use in surveys that model stylistic form.

Genetic Algorithms were used with good results to search the design space in the framework discussed in this dissertation. In future work, the performance of other optimization techniques, or indeed intelligent combinations of different optimization techniques, both stochastic and deterministic, could help to determine optimal ways to improve the probability and efficiency for finding the global minimum in design spaces that involve stylistic form and function goals. One such method is protocol-based multi-agent systems, which cooperatively uses multiple search protocols, which can each be adapted during optimization, to find higher quality solutions with greater efficiency (Landry & Cagan, 2011). Similarly, many promising methods have been developed to help manage multiple objectives in multi-objective optimization (Fonseca, 1995; Zitzler, 1999; Coello, 2000; Ignizio, 1976; Schniederjans, 1995). It may benefit future research to focus on incorporating other methods of intelligently managing the tradeoffs and synergies when optimizing design problems involving stylistic form and function. Another potential improvement that can be made to the optimization of stylistic form and function is to explore the benefits of intelligently finding the starting point or starting points for optimization algorithms, whether they are deterministic or stochastic. A number of methods have examined this issue (Brooks & Morgan, 1994) and may lend insight into how this can be refined for solving problems of stylistic form and function. One such possibility could be analogically inspired from the results of the experiment described in Section 6.2, where examples of existing vehicle designs, especially a wider range of them, led to improved management of form and functional goals for human designers. It may be possible that by using a wide range of existing vehicle design solutions, which may have more synergistic combinations of traits in their designs than a randomly generated population as starting points in optimization methods, or as the initial population for a Genetic Algorithm, could lead to higher quality solutions more efficiently.

Conjoint studies with logit models have been used to great effect for modeling stylistic form preference (Orsborn, 2009; Orsborn et al., 2010), allowing for quick and efficient mappings of utility functions for individual design parameters. Some limitations with the use of logit models are complications in modeling interactions between different design parameters (Brambor et al, 2005; Green, 1984; Hagerty, 1986). Accurate modeling of interaction effects is crucial in order to capture consumer judgments of stylistic form (Green & Srinivasan, 1990). Nevertheless, logit models are natively powerful at finding regions of values for individual parameters that are most promising for desired stylistic form goals without taking interactions into account. By coupling an Artificial Neural Network and a logit model, it is possible that a synergy could be created that uses the logit model to capture large general trends of individual parameters, and the ANN to capture nuances in interaction between features. Another possibility for combining these two methods is to use logit models to determine concise regions of interest for surveys and ANNs to model in greater detail, reducing wasted survey data points and computational expense for modeling uninteresting regions of the design space. One more possible synergy between logit models and ANNs involves the use of logit models to determine optimal parametric values to use either as a starting point for deterministic search algorithms, or to use logit models to determine the initial population for a genetic algorithm. This process may help to speed up searches for optimal designs and increase the probability of finding the global optimum by utilizing information from the logit model.

7.3.2 FUTURE COGNITIVE STUDIES

Additional research is needed to further confirm and explore the findings of the cognitive studies conducted in this research. The experiment discussed in Section 6.2 suggests that in addition to presenting non-visual inspirational information, the presentation of visual inspiration in the form of analogically relevant images enables participants to more effectively solve relevant design problems. A follow-up experiment should be conducted that examines how participants perform in the absence of any of the additional inspirational information. This condition would help to create a baseline to examine the effect of presenting non-visual inspirational information without visual inspiration. Similarly, this study should involve a larger number of participants than was used in the study discussed in Section 6.2 to better examine whether the significance of the statistically marginal finding that participants that received a wider range of inspirational material could better apply that information towards the design problem.

Following up to the experiment in Section 6.2, researchers should also examine in greater detail how and why the vehicle images enabled participants to solve the design problem more efficiently and effectively. This future research should examine whether the act of using visual aids helps to provide a deeper connection to the participant’s mental representation of the design problem, and if so, what other methods and forms of inspiration can benefit this representation. The effects of solution priming in the experiment should also be studied by using visual inspiration that does not contain solutions that prime the solution of the design problem.

In order to better understand how designers represent and manipulate different conflicting or synergistic goals in design problem solving, more experiments should be conducted in the future. The literature has suggested that human designers create multiple mental representations of the problem they are trying to solve. (Maher & Tang, 2003; Klahr & Dunbar, 1988) Research has also theorized about (Amabile, 1983; Crosby, 1968; Guetzkow, 1965) and studied (Shalley, 1991; Madjar & Shalley, 2008; Madjar & Oldham, 2006) the beneficial effects optimal goal switching in design problem solving, and warned against the effects of negatively enforced goal switching. Additional work needs to be done to study how designers natively switch between these goals and representations, how starting a participant on one representation or how enforcing a specific switching schedule might affect problem solving.

Preliminary work has been done by this researcher using the data and vehicle design interface discussed in Section 6.2 to study how participants natively switch between different goals of stylistic form and function. In the study that was conducted, participants were asked to design two vehicles to satisfy target goals of stylistic form and function while the vehicle design interface logged the path of design iterations that was traversed by each participant, and the values of the target goals. Sample plots of the data from Participant 18 solving the first design problem is shown in Figures 7.1 – 7.4. The goals for the first design problem was to design a vehicle that rates over 90% in sportiness, has a coefficient of drag below 0.25, and has over 75 cubic feet of volume. Figure 7.1 shows the sum of differences between the parameters of the vehicle design at each iteration and the final solution. Tracking this sum of differences shows the rate of convergence to the final design throughout problem solving. Figure 7.2 shows the sportiness rating of the vehicle design as determined by an ANN at each iteration. Figure 7.3 shows the aerodynamic coefficient of drag of the vehicle design at each iteration. Figure 7.4 shows the volume of the vehicle design at each iteration.

After studying the data trends from all eighteen participants, the results from this preliminary study suggest that designers focus most of their attention on one goal at a time, trying to bring that goal value toward the target. Once the goal value reaches the target, participants alter the design in an attempt to bring another goal toward its target without excessively sacrificing the first goal. This process iterates throughout the problem solving process until all goals are satisfied. This effect is illustrated in Figure 7.2, where the participant appears to be diligently focused on improving the sportiness rating of the design, resulting in a consistent rise in the sportiness rating, until the target rating of 90% sportiness is reached at around 140 iterations. After reaching this first goal, the participant appears to switch focus to increasing the volume of the design, as seen in Figure 7.4, and improving aerodynamics, as seen in Figure 7.3, while trying to maintain the sportiness rating at or near the target value.

From studying the data from this preliminary study, another trend was spotted, which is characterized in Figure 7.1. As seen in the literature (Kotovsky, Hayes, & Simon, 1985; Kotovsky & Simon, 1990), many participants start problem solving with an exploratory period where their designs do not appear to converge toward the final solution. This exploratory period is then followed by a period, or multiple periods, of consistent convergence toward the final solution. In the case of Participant 18, this exploratory period appears to have lasted for approximately 105 iterations, after which point the participant converged consistently toward the final solution. Some participants appear to converge for a period, followed by plateau, or flattening out, of the rate of convergence, followed by a large divergence, before the participant returns to a path of convergence toward the solution. It is hypothesized that during the initial exploratory period, participants do not have a clear idea of how to solve the problem. As the participant spends more time exploring, certain trends become more clear, and the participant begins to make steady progress based on those trends, converging toward the final solution. In cases where the rate of convergence plateaus, it is hypothesized that the participant had reached an impasse or local minimum in the design space, like the one that plagued a deterministic computational search method in Section 5.4.2.2. After a period of exploration to realize that there are no further improvements to be made from their current design, the participant makes bold changes, often causing divergence from the final solution, in an attempt to break free from the local minimum and find a successful solution.

In the second design problem, it was observed that participants needed a much shorter exploratory period on average before beginning to converge on solutions. It is hypothesized that this is caused because the problem representation, which used two of the same goals, was carried over into the second problem. It was shown in the experiment in Section 6.2 that participants in conditions where visual inspiration was presented required significantly fewer iterations to find a solution, and that presenting participants with a wider range of visual examples required marginally fewer iterations to find a solution. Due to the difficulty and lack of accuracy in identifying the exploratory and convergent periods from plots, it is currently unclear whether this faster problems solving was due to a shorter exploratory period from a better representation of the problem, from faster convergence in the convergent period, or both.

Additional research is needed to formalize the results of this preliminary study and to perform follow-up studies that more rigorously examine the hypotheses of the preliminary study. More specifically, this additional research should examine the salient circumstances in problem solving cause participants to switch goals, and what is discovered during the exploratory period that enables the participant to begin convergence on the solution. In concert with the experiments discussed in Section 6.1 and 6.2, this additional research could also examine what effect visual inspiration has on the length of time a participant spends exploring, and the rate of convergence after exploration has completed? It is possible that a more controlled experiment with fewer design choices and simpler goals may be beneficial to first examine these questions. Similarly, it is possible that an analogous heuristic or method for solving form and function design problems that mimics the convergence of human designers can be applied to computational design. One such possibility of this is the use of simulated annealing (Cagan & Kotovsky, 1997) in the solving of form and function design problems.

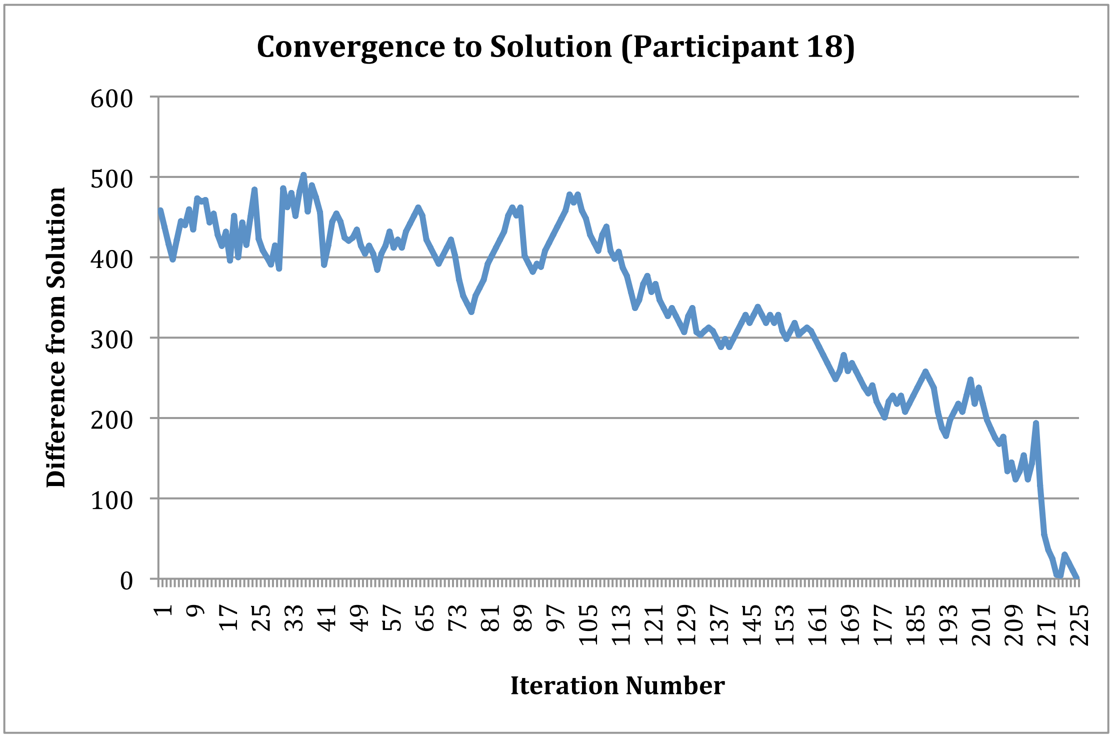

Figure 7.1 – Convergence to the solution for all iterations by Participant 18

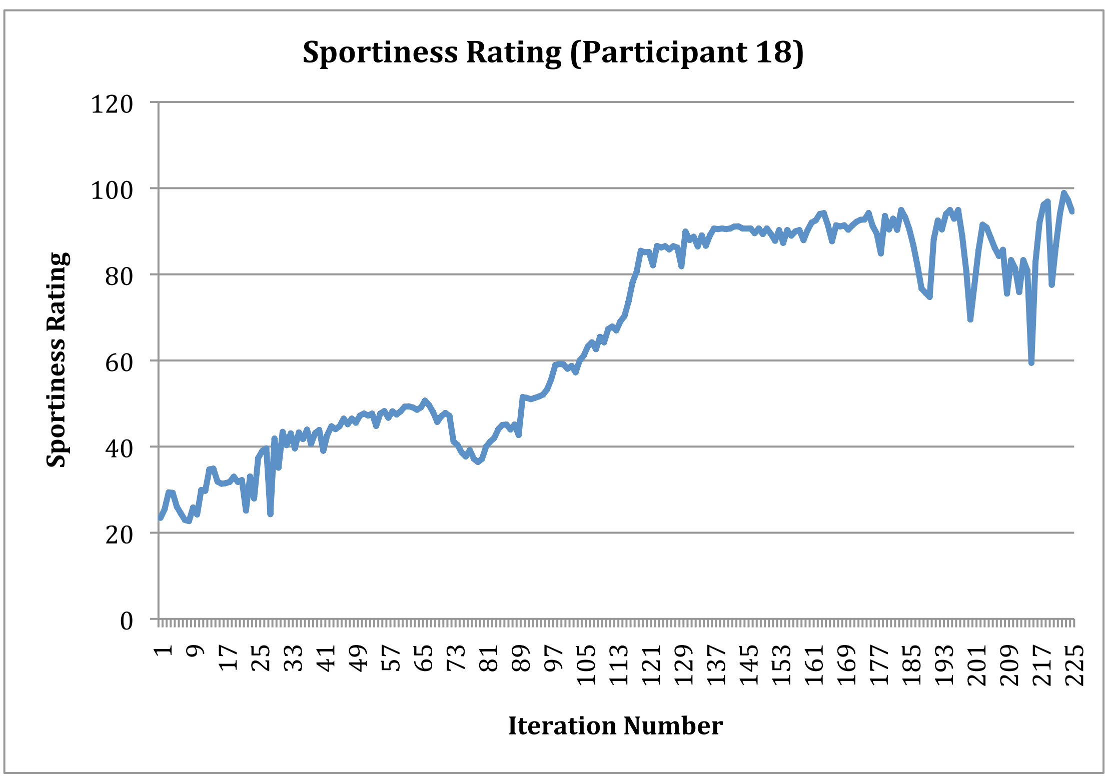

Figure 7.2 – Sportiness rating for all iterations by Participant 18

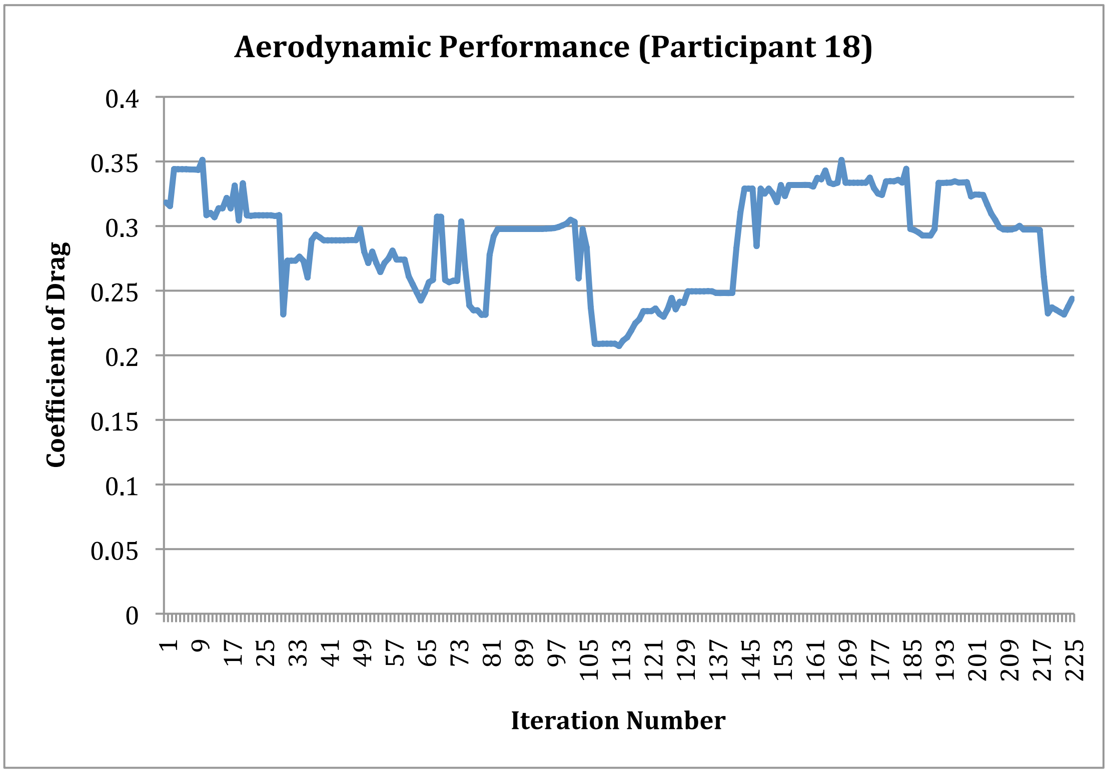

Figure 7.3 – Aerodynamic coefficient of drag for all iterations by Participant 18

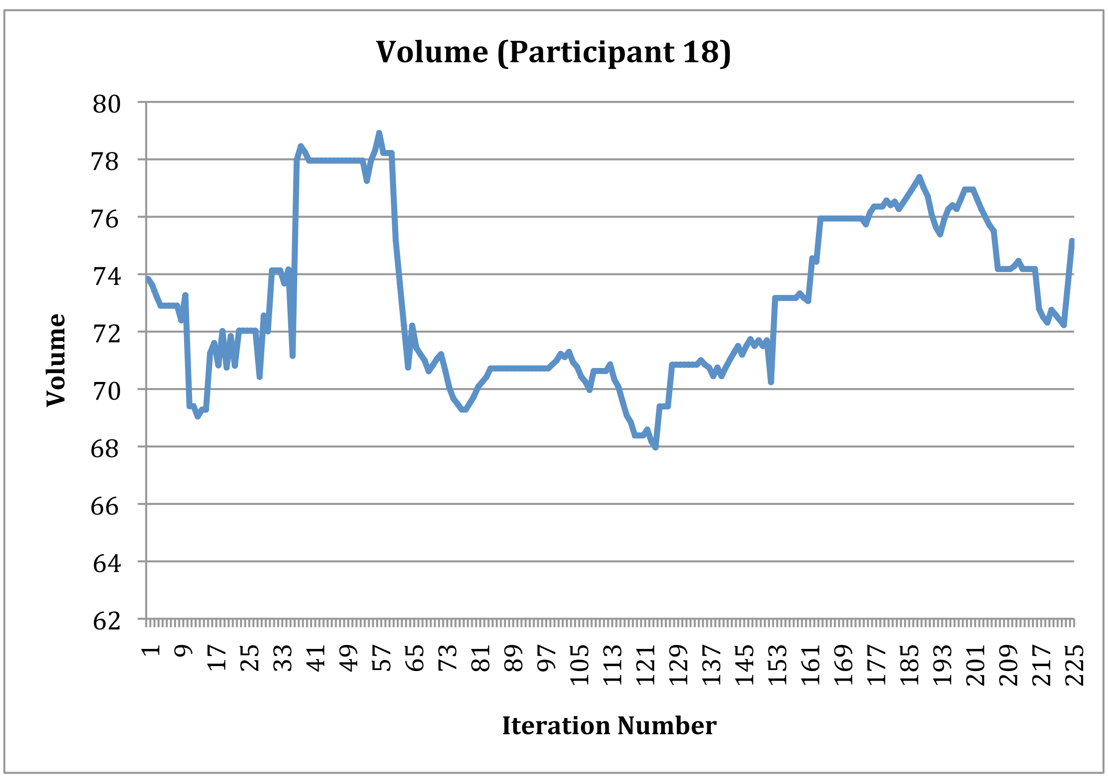

Figure 7.4 – Volume for all iterations by Participant 18

7.3.3 UNIFYING FUTURE COMPUTATIONAL AND COGNITIVE WORK

The goal of this dissertation is to examine how stylistic form and functional constraints can be brought together into a single cohesive framework. Humans and computers each have a unique set of abilities, and different ways to represent stylistic form and function. How can these native abilities be best utilized to maximize the efficiency and effectiveness of the design process? People have an innate ability to juggle goals of stylistic form and function. If this innate human method can be extracted and modeled, what can it tell us about how people design? Similarly, if this method is applied to manage a computational system that juggles stylistic form and function, will it inspire an improved method to computationally solve these problems? The inverse situation can also be true. Can an optimal representation-switching scheme on a computer inspire improved design methods that can yield more efficient human multitasking? Or is the maximum utility to be gained by combining the strengths of both humans and computers in a cooperative design environment?

A better understanding of how people represent and solve design problems, and a better understanding of how inspirational material can enrich these problem representations and aid problem solving could be combined to create a computational tool that monitors the progression of human designers and presents analogically appropriate inspirational material at optimal times throughout the problem solving process. This automatic inspiration system could help human designers of all skill levels to design to the best of their abilities to create truly innovative and optimal designs in a cooperative environment with computers. The reality of such a system is quite distant, but acts to motivate how each of the individual parts of research that is conducted in this dissertation and suggested in Sections 7.3.1 and 7.3.2 may eventually be assembled to accomplish amazing goals.

Future research should investigate the nuances of developing an optimal schedule for switching between multiple representations of stylistic form and function for both human designers and computational tools. Toward this goal, this future research should further the preliminary research discussed in Section 7.3.2 which aims to understand how human designers natively switch representations when given personal discretion, and how problem-solving performance is affected when other representation switching schedules are imposed. The optimal representation-switching schedule can also be determined for computational systems. A major goal of this research would be to investigate how optimal switching methods for human designers and computational systems can mutually inform each other to yield mutual improvements. In addition to enforcing switching schedules between these goals or representations, it may be possible to train a decision tree or a neural network to learn and mimic the switching behavior of human designers by presenting the machine learning technique with historical information of what designers do in certain situations in design problems. If successful, this method could lead to powerful heuristics that may help to improve the computational design process, and may lead to improved techniques that can be taught to human designers. Findings such as these would go beyond unifying stylistic form and function in design, and could lead to significant improvements in all elements of design.

7.4 CONCLUDING REMARKS

The work presented in this dissertation is in support of the thesis presented in Section 1.2.

> *A better understanding of how people behave as both designers and consumers can inspire designers to better manage product design goals and enable computers to treat form and functional considerations in a unified fashion using the same quantitative methods.*

A better understanding of how designers solve problems, and the problems associated with this problem solving, inspired two experiments that offer valuable insights that will help designers in the future to better manipulate product design goals for more efficient and effective management of synergies and tradeoffs in design. Also, a better understanding of the reasoning behind how consumer judgments are formed and manipulated inspired improved methods for surveying and modeling these consumer judgments. In turn, these surveys and modeling methods helped to enable computers to quantify stylistic form using machine-learning methods. Finally, these quantitative modeling methods opens the door for computers to treat form and functional considerations in a unified fashion using the same quantitative methods.

The method for learning stylistic form and generating new designs that satisfy both stylistic form and functional goals discussed in this dissertation has nearly limitless potential with a wide variety of product design categories, and the lessons learned about how designers can be inspired to better manage tradeoffs and synergies in design problems may one day help designers to better bridge the divide between stylistic form and function.

In short, the methods and experiments described in this dissertation have prescribed benefits toward the unification of stylistic form and function for both human designers and computational manipulation, and can be applied to the design of a wide range of products.

REFERENCES

[1] http://www.dpreview.com

[2] http://en.wikipedia.org/wiki/Volkswagen_A_platform

[3] http://dms.irb.hr/tutorial/tut_dtrees.php

Achiche, S., and Ahmed, S. (2008). "Mapping Shape Geometry and Emotions Using Fuzzy Logic". Proceedings from ASME 2008 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference.

Adeli, H., and Kumar, S. (1995). "Distributed genetic algorithm for structural optimization". Journal of Aerospace Engineeering, 8(3),156-163.

Agarwal, M., and Cagan, J. (1998). "A Blend of Different Tastes: The Language of Coffee Makers". Environment and Planning B: Planning and Design, Vol. 25, No. 2, 205-226.

Ai, C., Norton, E.C. (2003). "Interaction terms in logit and probit models". Economics Letters, 80, 123-129.

Aiken, L., and West, S. (1991). "Multiple Regression: Testing and Interpreting Interactions". London: Sage.

Amabile, T. M. (1983). "The social psychology of creativity". New York, Springer-Verlag.

Andersson, J. (2000). "A survey of multiobjective optimization in engineering design". Optimization, (LiTH-IKP-R-1097), 34.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C. and Qin, Y. L. (2004). "An integrated theory of the mind". Psychological Review, Vol. 111, No. 4, 1036-1060.

Antonsson, E. K., and Cagan, J. eds., (2001). "Formal Engineering Design Synthesis". Cambridge University Press, Cambridge, UK.

Baluja, S., Pomerleau, D., and Jochem, T. (1994). "Towards automated artificial evolution for computer-generated images". Connection Science, 6(2&3) 325-354.

Baxt, W.G. (1992), "Improving the Accuracy of an Artificial Neural Network Using Multiple Differently Trained Networks". Neural Computation, Vol. 4, No. 5, 772-780.

Becker, W., Yu, X., and Tu, J. (2007). “A generalized CFD learning and prediction system”. In: Proceedings of IEEE International Conference on Industrial Informatics, 2, 615-620.

Biles, J.A., Anderson, P.G., and Loggi, L.W. (1996). “Neural network fitness functions for a musical IGA”. In: The International ICSC Symposium on Intelligent Industrial Automation and Soft Computing.

Bishop, C. (1991). “Improving the Generalization Properties of Radial Basis Function Neural Networks”. Neural Computation, Vol. 3, No. 4, 579-588.

Blomkvist, O., Ekdahl, F., Gustafsson, A. (2007). “Non-geometric Plackett-Burman Designs in Conjoint Analysis”. In “Conjoint Measurement: Methods and Applications”. Fourth Edition, Eds. Gustafsson, A., Herrmann, A., Huber, F. Springer Berlin Heidelberg New York.

Bouchard, C., Lim, D., and Aoussat, A. (2003). “Development of a Kansei engineering system for industrial design: Identification of input data for KES”. 6th Asian Design International Conference, ADC, Tsukuba.

Box, G.E.P., and Meyer, R.D. (1993). “Finding the Active factors in fractionated screening experiments”. Journal of Quality Technology, 25, 94-105.

Brambor, T., Clark, W.R., Golder, M. (2005). “Understanding Interaction Models: Improving Emperical Analyses”. Political Analysis, 13, 1-20.

Breiman, L., Friedman, J., Olshen, R. A., and Stone, C. J. (1984). "Classification and regression trees". Wadsworth.

Brooks, S.P., and Morgan, B.J.T. (1994). “Automatic starting point selection for functional optimization”. Statistics and Computing, Vol. 4, No. 3, 173-177.

Bull, L. (1999). “On model-based evolutionary computation”. Soft Computing, 3, 76-82.

Cagan, J., and Kotovsky, K, (1997). “Simulated Annealing and the Generation of the Objective Function: A Model of Learning During Problem Solving”. Computational Intelligence, 13:4, 534-581.

Cagan, J., and Vogel, C. M., (2002). “Creating Breakthrough Products: Innovation from Product Planning to Program Approval”. Financial Times Prentice Hall, Upper Saddle River, NJ.

Cagdas, G. (1996). “A shape grammar: the language of traditional Turkish houses”. Environment and Planning B, 23(4), 443-464.

Calkins, D.E., Chan, W.T. (1998). “CDaero – A Parametric Aerodynamic Drag Prediction Tool”. SAE Technical Paper 980398, International Congress and Exposition, Feb. 1998, Detroit, MI, USA.

Campbell, D.T. (1960). “Blind variation and selective retentions in creative thought as in other knowledge processes”. Psychological Review, Vol. 67, 380-400.

Carr, G.W., Stapleford, W.R. (1981). “A Proposed Empirical Method for Predicting the Aerodynamic Drag of Cars”. Motor Industry Research Association (MIRA), England.

Casakin, H., and Goldschmidt, G. (1999). “Expertise and the Use of Visual Analogy: Implications for Design Education”. Des. Stud., Vol. 20, No. 2, 153–175.

Chan, W.T. (1997). “Empirical prediction of automobile drag coefficient”. MS Thesis, University of Washington.

Chang, C.C., and Shih, Y.Y. (2003). “A differential study on the product form perceptions of different age group users”. In: Proceedings of the 6th Asian Design International Conference

Chen, X., McKay, A., de Pennington, A., and Chau, HH. (2004). “Package shape design principles to support brand identity”. 14th IAPRI World Conference on Packaging, 14th IAPRI World Conference on Packaging.

Christensen, B.T. and Schunn, C.D. (2005) “Spontaneous access and analogical incubation effects”. Creativity Research Journal, Vol. 17, No. 2-3, 207-220.

Chrysikou, E. G., and Weisberg, R. W. (2005). “Following the wrong footsteps: Fixation effects of pictorial examples in a design problem-solving task”. Journal of Experimental Psychology-Learning Memory and Cognition, 31(5), 1134-1148.

Coello, C.A. (2000). “An updated survey of GA-based multiobjective optimization techniques”. ACM Computing Surveys, 32(2), 109-143.

Crosbie, A. (1968). “Creativity and performance in industrial organization”. London: Tavistock.

Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002). “A fast and elitist multiobjective genetic algorithm: NSGA-II”. IEEE Transactions on Evolutionary Computation, 6(2), 182-197.

DeMaris, A. (1990). "A Framework for the Interpretation of First-Order Interaction in Logit Modeling". Psychological Bulletin, Vol. 110, Issue 3, 557-570.

Dunbar, K. (1993). “Concept Discovery in a Scientific Domain”. Cognitive Science 17, 197-434.

Duncker, K. (1945). “On problem-solving”. Psychological Monographs, Vol. 58, No. 5 (Original work published in 1935).

Eckert, C.M., and Do, E.Y.L. (2006). “Special Issue: Understanding, representing, and reasoning about style”. Artificial Intelligence for Engineering Design, Analysis and Manufacturing, 20, 163-165.

Ewen, S. (1990). “All Consuming Images: The Politics of Style In Contemporary Culture”. Basic Books.

Fayyad, U.M., and Irani, K.B. (1992). “On the handling of continuous-valued attributes in decision tree generation”. Machine Learning, 8, 87-102.

Fayyad, U.M., and Irani, K.B. (1993). “Multi-Interval Discretization of Continuous-Valued Attributes for Classification Learning”. Proceedings of the 13th International Joint Conference on Artificial Intelligence. Morgan Kaufman, 1022-1029.

Fonseca, C. M., Fleming, P. J. (1995). “An overview of evolutionary algorithms in multiobjective optimization”. Evolutionary Computation, Vol. 3, No. 1, 1–16.

Forbus, K.D., Gentner, D. and Law, K. (1995). “MAC/FAC: a model of similarity-based retrieval”. Cognitive Science, Vol. 19, No. 2, 141-205.

Friedman, J.H. (1988). “Multivariate adaptive regression splines”. Technical report 102, Laboratory for Computational Statistics, Stanford University, Stanford, CA.

Friedrich, R. (1982). “In Defense of Multiplicative Terms in Multiple Regression Equations”. American Journal of Political Science 26:797–833.

Fu, K., Cagan, J., and Kotovsky, K., (2010). "Design Team Convergence: The Influence of Example Solution Quality". J. of Mech. Des., 132, 111005.

Fu, L., and Kara, L.B. (2011). “From engineering diagrams to engineering models: Visual recognition and applications”. Computer-Aided Design, Vol. 43, Issue 3, 278-292.

Fudenberg, D. and Tirole, J. (1983). “Game Theory”. MIT Press.

Fukushima, K. (1980). "Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position". Biological Cybernetics, 36[4], 193-202.

Fukushima, K. (1988). "Neocognitron: A hierarchical neural network capable of visual pattern recognition". Neural Networks, 1[2], 119-130.

Fukushima, K. (2003). "Neocognitron for handwritten digit recognition". Neurocomputing, 51, 161-180.

Gantovnik, V.B., Tiwari, S., Fadel, G.M., and Miao, Y. (2006). "Multi-objective vehicle layout optimization". 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference.

Gick, M.L. and Holyoak, K.J. (1980). "Analogical problem solving". Cognitive Psychology, Vol. 12, No. 3, 306-355.

Goicoechea, A., Hansen, D., and Duckstein, L. (1982). "Multiobjective Decision Analysis with Engineering and Business Applications". John Wiley and Sons.

Goldberg, D.E., (1989). "Genetic Algorithms in Search, Optimization, and Machine Learning". Addison Wesley.

Green, P.E. (1984). "Hybrid models for conjoint analysis: An expository review". Journal of Marketing Research, 21, 155-169.

Green, P.E., and Rao, V.R. (1971). "Conjoint Measurement for Quantifying Judgmental Data". Journal of Marketing Research, 8, (August), 355-363.

Green, P.E., and Srinivasan, V. (1978). "Conjoint analysis in consumer research: Issues and outlook". Journal of Consumer Research, 5, 103-123.

Green, P.E., and Srinivasan, V. (1990). "Conjoint analysis in marketing: New developments with implications for research and practice". Journal of Marketing, 17, 3-19.

Griffith, N., and Todd, P.M., eds. (1999). "Musical Networks: Parallel Distributed Perception and Performance". Cambridge, MA: MIT Press/Bradford Books.

Grignon, P.M., and Fadel, G.M. (2004). "A GA based configuration design optimization method". Journal of Mechanical Design, 126, 6-15.

Guan, L. (1995). "Feature based aerodynamic drag coefficient metamodel". MS Thesis, University of Washington.

Guetzkow, H. (1965). "The creative person in organizations". In G.A. Steiner (Ed.), The creative organization, University of Chicago Press, 35-49.

Hagerty, M.R. (1986). "The cost of simplifying preference models". Marketing Science, 5, 198-319.

Hamada, M., Wu, C.F.J. (1992). “Analysis of designed experiments with complex aliasing”. Journal of Quality Technology, 24, 130-137.

Hans, A. E., (1988). “Multicriteria optimization for highly accurate systems”. in Multicriteria Optimization in Engineering and Sciences, 309–352.

Harrison, L. E., and Huntington, S. P. eds. (2001). “Culture Matters: How Values Shape Human Progress”. Basic Books.

Hashem, S., and Schmeiser, B. (1995). “Improving model accuracy using optimal linear combinations of trained neural networks”. IEEE Trans. on Neural Networks, Volume 6, Issue 3, 792-794.

Hazelrigg, G. A. (1998) “A Framework for Decision-Based Engineering Design”. Journal of Mechanical Design, Vol. 120, Issue 4, 653-659.

Holyoak, K.J. and Koh, K. (1987). “Surface and structural similarity in analogical transfer”. Memory and Cognition, Vol. 15, No. 4, 332-340.

Horn, J., Nafpliotis, N., and Goldberg, D.E. (1994). “A niched Pareto genetic algorithm for multiobjective optimization”. In: Proceedings of the First IEEE Conference on Evolutionary Computation, 1, 82-87.

Hsiao, S.W., and Huang, H.C. (2002). “A neural network based approach for product form design”. Design Studies, 23(1), 67-84.

Hsiao, S.W., and Liu, M.C. (2002). “A morphing method for shape generation and image prediction in product design”. Design Studies, 23(5), 533-556.

Ignizio, J. (1976). “Goal Programming and Extensions”. Lexington Books.

Jansson, D. G., and Smith, S. M. (1991). “Design fixation”. Design Studies, 12(1), 3-11.

Jin, Y., Levitt, R. (1996) “The Virtual Design Team: A Computational Model of Project Organizations”. Computational and Mathematical Organization Theory, 2:3 171-196.

Johnson, R. (1974). “Trade-off analysis of consumer values”. Journal of Marketing Research, 11, 121-127.

Jolliffe, I.T. (2002). “Principal Component Analysis”. 2nd ed., Springer.

Kam, C., Franzese, R. (2003). “Modeling and interpreting interactive hypotheses in regression analysis: A brief refresher and some practical advice”. Unpublished manuscript, University of Michigan.

Kim, H.S., and Cho, S.B. (2000). “Application of interactive genetic algorithm to fashion design”. Engineering Applications of Artificial Intelligence, 13(6), 635-644.

Klahr, D., and Dunbar, K. (1988). “Dual space search during scientific reasoning”. Cognitive Science, 12, 1-48.

Kolodner, J. L., 1997, “Educational Implications of Analogy: A View from Case-Based Reasoning”. Am. Psychol., 52 1 , 57–66.

Kotovsky, K., Hayes, J.R., Simon, H.A. (1985). “Why are some problems hard? Evidence from Tower of Hanoi”. Cognitive Psychology, 17, 248-294.

Kotovsky, K., Simon, H.A. (1990). “What makes some problems really hard: Explorations in the problem space of difficulty”. Cognitive Psychology, 22, 143-183.

Landry, L.H., and Cagan, J. (2011). “Protocol-Based Multi-Agent Systems: Examining the Effect of Diversity, Dynamism, and Cooperation in Heuristic Optimization Approaches”. J. Mech. Des., Vol. 133, Issue 2.

Lawrence, S., Giles, C.L., Tsoi, A.C., and Back, A.D. (1997). “Face recognition: a convolutional neural-network approach”. IEEE Transactions on Neural Networks, Vol. 8, Issue 1, 98-113.

Lee, H.C., and Tang, M.X. (2009). “Evolving product form designs using parametric shape grammars integrated with genetic programming”. Artificial Intelligence for Engineering Design, Analysis and Manufacturing, 23, 131-158.

Levenberg, K. (1944). "A Method for the Solution of Certain Non-Linear Problems in Least Squares". The Quarterly of Applied Mathematics 2: 164–168.

Li, H., Azarm, S. (2000). “Product design selection under uncertainty and with competitive advantage”. Journal of Mechanical Design, Vol. 122, Issue 4, 411-418.

Linsey, J.S., Green, M.G., Murphy, J.T. and Wood, K.L. (2005). “Collaborating to success: an experimental study of group idea generation techniques”. In Proceedings of the DETC’98, 2005 ASME Design Engineering Technical Conferences, Long Beach, CA.

Linsey, J.S., Tseng, I., Fu, K., Cagan, J., Wood, K. L., and Schunn, C. D., 2010, “A Study of Design Fixation, Its Mitigation and Perception in Design Faculty”. J. of Mech. Des., 132, 041003.

Linsey, J.S., Wood, K.L. and Markman, A.B. (2008). “Modality and Representation in Analogy”. Artificial Intelligence in Engineering Design Analysis and Manufacturing ,Vol. 22, No. 2.

Lophaven, S., Nielsen, H. B., Sondergaard, J. (2002). “DACE A MATLAB Kriging Toolbox, Informatics and Mathematical Modelling”. Technical University of Denmark.

Luce, R.D. & Tukey, J.W. (1964). “Simultaneous conjoint measurement: a new scale type of fundamental measurement”. Journal of Mathematical Psychology, 1, 1–27.

Madjar, N., and Oldham, G. R. (2006). “Task rotation and polychronicity: Effects on individuals’ creativity”. Human Performance, 19: 117-131.

Madjar, N., Shalley, C. E. (2008). “Multiple Tasks’ and Multiple Goals’ Effect on Creativity: Forced Incubation or Just a Distraction?” Journal of Management, 34, 786-805.

Maher, M.L., and Tang, H.H. (2003). “Co-evolution as a computational and cognitive model of design”. Research in Engineering Design, 14, 47-63.

Marler, R. T., Arora, J. S. (2009). “The weighted sum method for multi-objective optimization: new insights”. Struct Multidisc Optim, Springer.

Marler, R. T., Arora, J. S. (2004). “Survey of multi-objective optimization methods for engineering”. Struct Multidiscipl Optim 26:369-395

Marsh, R. L., Ward, T. B., and Landau, J. D., 1999, “The Inadvertent Use of Prior Knowledge in a Generative Cognitive Task”. Mem. Cognit., 27 1, 94-105.

Marquardt, D. W., (1963). “An Algorithm for Least-Squares Estimation of Nonlinear Parameters”. Journal Society Industrial Applied Mathematics, Vol. 11, No. 2, June 1963, 431-441.

McCormack, J. P., Cagan, J., and Vogel, C. M. (2004). “Speaking the Buick Language: Capturing, Understanding, and Exploring Brand Identity with Shape Grammars”. Design Studies, Vol. 25, 1-29.

Michalek, J.J., F.M. Feinberg and P.Y. Papalambros (2004) “An optimal marketing and engineering design model for product development using analytical target cascading”. Proceedings of the TMCE Conference, April 12-16, Lausanne, Switzerland.

Michalek, J.J., O. Ceryan, P.Y. Papalambros, and Y. Koren (2006) “Balancing marketing and manufacturing objectives in product line design”. ASME Journal of Mechanical Design, Vol. 128, No. 6, 1196-1204.

Mitchell, T. (1997). “Machine Learning”. McGraw Hill.

Montazeri-Gh, M. and Poursamad, A. (2006). “Application of genetic algorithm for simultaneous optimisation of HEV component sizing and control strategy”. Int. J. Alternative Propulsion, Vol. 1, No. 1,63–78.

Morgan, G. (1998). “Images of Organization: The Executive Edition [Abridged]”, Berrett-Koehler Publishers.

Moss, J., Kotovsky, K., and Cagan, J. (2007). “The influence of open goals on the acquisition of problem relevant information”. Journal of Experimental Psychology: Learning, Memory, and Cognition, 33(5), 876-891.

Moss, J., Kotovsky, K., and Cagan, J. (2008). “Understanding the role of open goals in problem solving: Impasses, fixation, and hints”. Manuscript submitted for publication.

Nadler, D., Tushman, M., Nadler, M. (1997). “Competing by Design: The Power of Organizational Architecture”. Second Edition, Oxford University Press.

Nagamachi, M. (2002). “Kansei engineering as a powerful consumer-oriented technology for product development”. Applied Ergonomics, 33, 289-294.

Norton, E.C., Wang, H., Ai, C. (2004). "Computing interaction effects and standard error in logit and probit models". The Stata Journal, Vol. 4, No. 2, 154-167.

Nussbaum, B. (2002) Foreword. In Cagan, J., and Vogel, C. M., (2002). “Creating Breakthrough Products: Innovation from Product Planning to Program Approval”. Financial Times Prentice Hall, Upper Saddle River, NJ.

Olhofer, M., Jin, Y., Sendhoff, B. (2001). “Adaptive encoding for aerodynamic shape optimization using Evolution Strategies”. In: Proceedings of IEEE Congress on Evolutionary Computation, 1, 576-583.

Orsborn, S. (2007). “Quantifying aesthetic preference through statistics applied to an agent-based shape grammar implementation”. PhD Dissertation, Carnegie Mellon University.

Orsborn, S., and Cagan, J. (2009a)."Automatically Generating Form Concepts According to Consumer Preference: A Shape Grammar Implementation with Software Agents". ASME Journal of Mechanical Design, Vol. 131, 121007-1-10.

Orsborn, S., Cagan, J. (2009b). "Multiagent Shape Grammar Implementation: Automatically Generating Form Concepts According to a Preference Function". Journal of Mechanical Design, Vol. 131, 121007-1-10.

Orsborn, S., Cagan, J., Boatwright, P. (2009). “Quantifying Aesthetic Form Preference in a Utility Function”. Journal of Mechanical Design, Vol. 131, 061001-1-10.

Orsborn, S., Cagan, J., Pawlicki, R., and Smith, R.C. (2006). “Creating cross-over vehicles: Defining and combining vehicle classes using shape grammars”. Artificial Intelligence for Engineering Design, Analysis and Manufacturing, 20, 217-246.

Osborn, A. (1957). “Applied imagination”. Scribner, New York, NY.

Osgood, C., Suci, G. and Tannenbaum, P. (1957). “The Measurement of Meaning”. University of Illinois Press, Urbana, Il.

Parsons, G., Carlson, A. (2009). “Functional Beauty”. Oxford University Press, USA.

Perttula, M. K., and Liikkanen, L. A. (2006). “Structural tendencies and exposure effects in design idea generation”. Paper presented at the ASME 2006 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, Philadelphia, PA.

Prudêncio, R. B. C., and Ludermir, T. B. (2002). “Neural Network Hybrid Learning: Genetic Algorithms and Levenberg-Marquadt”. In: 26th Annual Conference of the Gesellschaft fuer Klassifikation.

Pugliese, M. and Cagan, J. (2002). "Capturing a Rebel: Modeling the Harley-Davidson Brand through a Motorcycle Shape Grammar". Research in Engineering Design, Vol. 13, No. 3, 139-156.

Purcell, A. T., and Gero, J. S. (1996). “Design and other types of fixation”. Design Studies, 17, 363-383.

Quinlan, J.R. (1990). “Decision trees and decisionmaking”. IEEE Transactions on Systems, Man, and Cybernetics, 20(2), 339-346.

Quinlan, J.R. (2006). “Learning with continuous classes”. In Proceedings of AI’92, Singapore: World Scientific, 1992, 343-348.

Reid, T., Gonzalez, R., Papalambros, P. (2009). “A methodology for quantifying the perceived environmental friendliness of vehicle silhouettes in engineering design”. In: Proceedings of the ASME 2009 International Design Engінerring Technical Conferences and Computers and Information in Engineering Conference.

Schniederjans, M. (1995). “Goal programming: methodology and applications”. Kluwer Academic Publishers.

Shah, J.J. (1998). “Experimental investigation of progressive idea generation techniques in engineering design”. In: Proceedings of the DETC’98, 1998 ASME Design Engineering Technical Conferences, Atlanta, GA.

Shalley, C. E. (1991). “Effects of productivity goals, creativity goals, and personal discretion on individual creativity”. Journal of Applied Psychology, 76: 179-185.

Shiau, C.S., Tseng, I.H., Heutchy, A.W., and Michalek, J. (2007). "Design optimization of a laptop computer using aggregated and mixed logit demand models with consumer survey data". Proceedings of the ASME International Design Engineering Technical Conferences, Las Vegas, Nevada, USA.

Shrout, P., Fleiss, J. (1979). “Intraclass Correlations: Uses in Assessing Rater Reliability”. Psychological Bulletin, Vol. 86, No. 2, 420-428.

Simonton, D.K. (1999). “Creativity as blind variation and selective retention: is the creative process Darwinian?” Psychological Inquiry, Vol. 10, No. 4, 309-328.

Simpson, T.W. (2004). “Product platform design and customization: Status and promise”. Artificial Intelligence for Engineering Design, Analysis and Manufacturing, 18, 3-20.

Simard, P.Y., Steinkraus, D., Platt, J.C. (2003). “Best Practices for Convolutional Neural Networks Applied to Visual Document Analysis”. Proceedings of the Seventh International Conference on Document Analysis and Recognition.

Smith, R.C., Pawlicki, R., Kokai, I., Finger, J., and Vetter, T. (2007). “Navigating in a shape space of registered models”. In: IEEE Transactions on Visualization and Computer Graphics, 13(6), 1552-1559.

Smith, S.M. and Blankenship, S.E. (1991). “Incubation and the persistence of fixation in problem solving”. American Journal of Psychology, Vol. 104, 61-87.

Stacey, M. (2006). “Psychological challenges for the analysis of style”. Artificial Intelligence for Engineering Design, Analysis and Manufacturing, 20, 167-184.

Steuer, R. (1989). “Multiple Criteria Optimization: Theory, Computation, and Application”. Robert E. Krieger Publishing.

Stiny, G. (1980). “Introduction to shape and shape grammars”. Environment and Planning B, 7(3), 343-351.

Stiny, G., Mitchell, W. J., (1978). “The Palladian grammar”. Environment and Planning B, 5(1), 5-18.

Szykman, S., and Cagan, J. (1995). “A Simulated Annealing-Based Approach to Three-Dimensional Component Packing”. ASME Journal of Mechanical Design, Vol. 117, No. 2(A), 308-314.

Takagi, H. (2001). “Interactive evolutionary computation: Fusion of the capabilities of EC optimization and human evaluation”. In: Proceedings of the IEEE, 89(9), 1275-1296.

Train, K. (2003). “Discrete Choice Methods with Simulation”. Cambridge University Press, New York.

Twomey, J.M., Smith, A.E. (2002). “Bias and variance of validation methods for function approximation neural networks under conditions of sparse data”. Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on, Vol. 28, Issue 3, 417-430.

Tseng, I., J. Moss, J. Cagan, K. Kotovsky, (2008). “The role of timing and analogical similarity in the stimulation of idea generation in design”. Design Studies, Vol. 29, 203-221.

Tsutsumi, K., Sasaki, K. (2008). “Study on shape creation of building’s roof by evaluating aesthetic sensibility”. Mathematics and Computers in Simulation, Vol. 77 Issue 5-6.

Ullman, D. G. (2003). “The Mechanical Design Process”. Third Edition, McGraw Hill.

Ward, T. B., (1994). “Structured Imagination: The Role of Category Structure in Exemplar Generation”. Appl. Cognit. Psychol., 27, 1, 1–40.

Wassenaar, H. J., Chen, W. (2003). “An approach to decision-based design”. In: Proceedings of the ASME 2001 Design Engineering Technical Conferences and Computers and Information in Engineering Conference.

Wassenaar, H. J., Chen, W., Cheng, J. and Sudjianto, A. (2005). “Enhancing Discrete Choice Demand Modeling for Decision-Based Design”. ASME J. Mech. Des., 127(4), 514-523.

Wiley, J. (1998). “Expertise as mental set: the effects of domain knowledge in creative problem solving”. Memory and Cognition, Vol. 26, No. 4, 716-730.

Wittink, D.R., Vriens, M., Burhenne, W. (1994). “Commercial Use of Conjoint Analysis in Europe: Results and Critical Reflections”. International Journal of Research in Marketing, 11, 41-52.

Wood, K. L., Greer, J. L. (2001). “Function-based Synthesis Methods in Engineering Design”. Chapter Six in Antonsson, E. K., and Cagan, J. eds., (2001). “Formal Engineering Design Synthesis”, Cambridge University Press, Cambridge, UK.

Yanagisawa, H., and Fukuda, S. (2004). “Development of interactive industrial design support system considering customer’s evaluation”. JSME International Journal, 47(2), 762-769.

Yanagisawa, H., and Murakami, T. (2008). “Factors affecting viewpoint shifts when evaluating shape aesthetics towards extracting customer’s latent needs of emotional quality”. In: Proceedings of the ASME 2008 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference.

Yannou, B., Dihlmann, M., Awedikian, Roy. (2008). “Evolutive Design of Car Silhouettes”. Proceedings from ASME 2008 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference.

Yasuda, K., Furuta, H., Kobayashi, T. (1995). “Aesthetic design system of structures using neural network and image database”. Proceedings of the 3rd International Symposium on Uncertainty Modelling and Analysis.

Yu, J., Goos, P., Vandebroek, M. (2007). "The importance of attribute interactions in conjoint choice design and modelling". Katholieke Universiteit Leuven, Working Paper KBI 0601.

Zadeh, L. (1963). “Optimality and non-scalar-valued performance criteria”. IEEE Trans- actions on Automatic Control, Vol. 8, 59–60.

Zhang, G.P. (2000). “Neural networks for classification: A survey”. IEEE Trans. on Systems, Man, and Cybernetics, Part C: Applications and Reviews, Vol. 30, Issue 4, 451-462.

Zitzler, E., Thiele, L. (1999). “Multiobjective evolutionary algorithms: A comparative case study and the strength pareto approach”. IEEE Transactions on Evolutionary Computation, Vol. 3, No. 4, 257–271.

APPENDIX

A MATLAB PARAMETRIC CAR MODEL CODE

```
function Funct_Drawcar(chromosome)

chassislength = 775;
beltangle = chromosome(1)*0.012;
noserake = chromosome(2)*.65;
groundclearance = chromosome(3)*0.40+20;
trimheight = chromosome(4)*0.35+35;
cowlheight = chromosome(4)*0.30+40;
roofheight = chromosome(5)*0.30+40;
hoodlength = chromosome(6)+125;
trunklength = chromosome(7)*0.95+5;
frontscreenrake = chromosome(8)*0.25+40;
rearscreenrake = chromosome(9)*0.45+20;
wheelsize = chromosome(10)*0.15+45;
cfx = chromosome(11)*0.50+175;
crx = chromosome(12)*0.50+650;

%%floorpan
    c1P0x = chassislength;
    c1P0y = 100+groundclearance;
    c1P1x = chassislength-50;
    c1P1y = 100+groundclearance-15;
    c1P2x = 110;
    c1P2y = 100+groundclearance-15;
    c1P3x = 100;
    c1P3y = 100+groundclearance;

    %%frontbumper
    c2P0x = c1P3x;
    c2P0y = c1P3y;
    c2P1x = c1P3x-5;
    c2P1y = c1P3y+8;
    c2P2x = c1P3x-7;
    c2P2y = c1P3y+trimheight-3;
    c2P3x = c1P3x;
    c2P3y = c1P3y+trimheight;

    %%hood
    c3P0x = c2P3x;
    c3P0y = c2P3y;
    c3P1x = c2P3x+noserake*hoodlength/200;
    c3P1y = c2P3y+cowlheight;
    c3P2x = c2P3x+(noserake*hoodlength*.01);
    c3P2y = c2P3y+(cowlheight*0.8)-beltangle/hoodlength;
    c3P3x = c2P3x+hoodlength;
    c3P3y = c2P3y+cowlheight;

    %%windshield
    c4P0x = c3P3x;
    c4P0y = c3P3y;
    c4P1x = c3P3x;
```

```
c4P1y = c3P3y;
c4P2x = c3P3x+(roofheight*tan(frontscreenrake/57));
c4P2y = c3P3y+roofheight;
c4P3x = c3P3x+(roofheight*tan(frontscreenrake/57));
c4P3y = c3P3y+roofheight;

%%rearbumper
c8P0x = c1P0x+5;
c8P0y = c1P0y + trimheight + (chassislength*sin(beltangle/57));
c8P1x = c1P0x+15;
c8P1y = c1P0y + trimheight + (chassislength*sin(beltangle/57)) -3;
c8P2x = c1P0x+5;
c8P2y = c1P0y+5;
c8P3x = c1P0x;
c8P3y = c1P0y;

%%trunk
c7P0x = c8P0x-trunklength;
c7P0y = c8P0y + cowlheight + (chassislength*sin(beltangle/57));
c7P1x = c8P0x-trunklength/8;
c7P1y = c8P0y + cowlheight + (chassislength*sin(beltangle/57));
c7P2x = c8P0x;
c7P2y = c8P0y + cowlheight;
c7P3x = c8P0x;
c7P3y = c8P0y;

%%rearscreen
c6P0x = c7P0x-(roofheight*tan(rearscreenrake/57));
c6P0y = c7P0y-(chassislength*sin(beltangle/57))+roofheight;
c6P1x = c7P0x-(roofheight*tan(rearscreenrake/57)*.6);
c6P1y = c7P0y-(chassislength*sin(beltangle/57))+roofheight*.7;
c6P2x = c7P0x;
c6P2y = c7P0y;
c6P3x = c7P0x;
c6P3y = c7P0y;

%%roof
c5P0x = c4P3x;
c5P0y = c4P3y;
c5P1x = c4P3x+(c6P0x-c4P3x)*.4;
c5P1y = c4P3y+roofheight*.5;
c5P2x = c6P0x-(roofheight*tan(rearscreenrake/57)*.4);
c5P2y = c6P0y+roofheight*.2;
c5P3x = c6P0x;
c5P3y = c6P0y;

cfy=100+wheelsize;

cry=100+wheelsize;

c1x = [c1P0x,c1P1x,c1P2x,c1P3x];
c1y = [c1P0y,c1P1y,c1P2y,c1P3y];
c2x = [c2P0x,c2P1x,c2P2x,c2P3x];
c2y = [c2P0y,c2P1y,c2P2y,c2P3y];
c3x = [c3P0x,c3P1x,c3P2x,c3P3x];
c3y = [c3P0y,c3P1y,c3P2y,c3P3y];
c4x = [c4P0x,c4P1x,c4P2x,c4P3x];
c4y = [c4P0y,c4P1y,c4P2y,c4P3y];
c5x = [c5P0x,c5P1x,c5P2x,c5P3x];
c5y = [c5P0y,c5P1y,c5P2y,c5P3y];
c6x = [c6P0x,c6P1x,c6P2x,c6P3x];
```

```
    c6y = [c6P0y,c6P1y,c6P2y,c6P3y];
    c7x = [c7P0x,c7P1x,c7P2x,c7P3x];
    c7y = [c7P0y,c7P1y,c7P2y,c7P3y];
    c8x = [c8P0x,c8P1x,c8P2x,c8P3x];
    c8y = [c8P0y,c8P1y,c8P2y,c8P3y];

    Funct_Bezier(c1x,c1y,100);
    Funct_Bezier(c2x,c2y,100);
    Funct_Bezier(c3x,c3y,100);
    Funct_Bezier(c4x,c4y,100);
    Funct_Bezier(c5x,c5y,100);
    Funct_Bezier(c6x,c6y,100);
    Funct_Bezier(c7x,c7y,100);
    Funct_Bezier(c8x,c8y,100);

    circle([cfx,cfy],wheelsize,100,'LineWidth');
    circle([crx,cry],wheelsize,100,'LineWidth');

    plot(0,0,'LineWidth',0.0001);
    plot(900,500,'LineWidth',0.0001);

    axis equal;

end
```

B INSPIRATIONAL MATERIAL FOR SECTION 6.1

Three-clocks description

Please read and study this information until we tell you to stop. You have three minutes. There will be a quiz on this material at the end.

GRANDFATHER CLOCK

A grandfather clock uses the constant period of a swinging pendulum to provide a continuous and stable reference frequency. This pendulum in turn drives the escapement, which is generally a gear and a pair of stops, which are actuated by the pendulum, that allow one tooth of the escapement's gear to "escape" after each full swing of the pendulum. The engagement of the two stops results in the characteristic "tick" and "tock" sounds of a clock. The escapement's gear is connected to a series of gears that control the relative speed of rotation between the escapement and the hands of the clock, the bells, and other elements of the clock. The energy to drive the hands is provided by a set of dropping weights that drop a small amount per cycle. These weights also provide just enough energy to the pendulum to overcome friction via the escapement.

WINDUP CLOCK

A windup clock uses the constant period of a spring powered rotating mass or flywheel, which works much like a pendulum in providing a continuous and stable reference frequency. This flywheel drives an escapement much like as used in a grandfather clock, which in turn drives the hands and other functions of the clock. The flywheel is generally small and turns at a much higher frequency than a pendulum, which results in the ability to drive a second hand. The power to drive the flywheel and the hands is provided by a spring, which is tensioned by winding.

QUARTZ WRISTWATCH

A quartz wristwatch uses an electronic quartz crystal oscillator to provide a constant period. Most battery-powered crystal clocks use a 32.768 kHz oscillator. Using the piezoelectric effect, an excited crystal generates voltage pulses, which are then divided down using a frequency divider or counter and used to drive a tiny electric motor, which in turn drives the hands and other functions of the wristwatch.

Three distant devices description:

Please read and study this information until we tell you to stop. You have three minutes. There will be a quiz on this material at the end.

HEART RATE MONITOR

A heart rate monitor is a device that allows a user to measure his or her heart rate in real time. It usually consists of two elements: a chest strap transmitter and a wrist receiver (which usually doubles as a watch). Strapless heart rate monitors are available as well, but lack some of the functionality of the original design. Advanced models additionally measure heart rate variability to assess a user's fitness.

The chest strap has electrodes in contact with the skin to monitor the electrical voltages in the heart. When a heartbeat is detected a radio signal is sent out which the receiver uses to determine the current heart rate.

CASSETTE TAPE DECK

A tape recorder, tape deck, reel-to-reel tape deck, cassette deck or tape machine is an audio storage device that records and plays back sound using magnetic tape, either wound on a reel or in a cassette, for storage. It records a fluctuating signal by moving the tape across a tape head that polarizes the magnetic domains in the tape in proportion to the audio signal.

Professional recorders usually use a simple three-motor scheme. One motor with a constant rotation speed provides traction for the leading wheel that is usually combined with a capstan and flywheel to ensure that the tape speed does not fluctuate. The other two motors apply constant torque to maintain the tape's tension or wind the tape quickly. Cheaper models use a single motor for all required functions. There are also variants with two motors, in which one motor is used for rewinding only.

WATER METER

A water meter is a device used to measure water usage. Water meters are normally used at every residence and commercial building in a public water system. Water meters can also be used at the water source, well, or throughout a water system to determine flow through that portion of the system. Water meters typically measure and display total usage in US gallons, cubic feet, or cubic meters on a mechanical or electronic register.

Water meters typically fall into two categories. A displacement type water meters relies on the water to physically displace the moving measuring element in direct relation to the amount of water that passes through the meter. The piston or disk moves a magnet that drives the register. A velocity type water meter measures the velocity of flow though a meter of a known internal capacity. The speed of the flow can then be converted into volume of flow for usage.