Exploring Design Methods for Dynamic User Preferences

by

Mojtaba Arezoomand

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mechanical Engineering)
in the University of Michigan
2021

Doctoral Committee:

Professor Jesse Austin-Breneman, Chair
Professor Shanna Daly
Professor Neda Masoud
Professor Elliott Rouse

Mojtaba Arezoomand
mojiar@umich.edu
ORCID iD: 0000-0001-6253-5666

ACKNOWLEDGMENTS

I am extremely thankful of everyone who has helped me with mentorship, advice, and support throughout my Ph.D. journey. To my advisor Jesse Austin-Breneman, thank you for your continuous support and guidance, and thank you for being as much a friend as an academic supervisor. Maintaining good relationship with graduate students is hard on its own, being a friend as well makes it even harder. Jesse, you were both an excellent advisor and a great friend. Thank you to my committee members, Shanna Daly, Neda Masoud, and Elliott Rouse for their invaluable insights, guidance, and thoughtful feedback that made my work more meaningful and impactful.

Many thanks to my amazing lab mates Suzanne Chou, John Meluso, and Meira Chefitz for being wonderful people as well as great friends. I consider myself extremely lucky to get to know you and have the chance to spend time with you and work along side you. Also thank you to Claudia Zimmerman and Jack Campau for your help with my last study, I would not have been able to finish my work on time without your help.

Finally, thank you to my parents, and family who have made as much sacrifice as I did over the past five years and made it possible for me to finish my degree here at the University of Michigan. I owe all my past and future accomplishments in life to you, and your continuous support and unconditional love. I hope to see you very soon.

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE

LIST OF TABLES

TABLE

ABSTRACT

New product development requires engineering designers to translate abstract user needs into concrete engineering specifications. Formal methods for engineering design offer strategies to guide decision-making throughout this process. Many methods assume that the artifact will be mass-manufactured for sale within a competitive market environment. The objective of these methods is to define product attributes which will maximize user preference relative to other options and therefore improve the likelihood of purchase. However, the emergence of the Internet of Things, smart products, and highly-connected systems has led to a growth in products sold as a service. Thus, user preferences during the entire product lifetime becomes increasingly important. This dissertation will extend design research to investigate approaches which maximize user preferences over the lifetime of the product. The first study examines the coupled interaction between product attributes and contextual parameters in large connected systems. The second study explores the use of reinforcement learning algorithms to design adaptive hardware systems which respond to dynamic user preferences. The third study explores the relation between reinforcement learning algorithms' behavior with user willingness for providing feedback in adaptive hardware systems. Together these studies create a road map towards prescriptive design methods which enable engineers to design systems which tailor their functionality to changing user needs over the lifetime of the product.

The first study focuses on the design of large connected user-product-environment systems. Usage context and more specifically, the dynamic usage variables that are impacted indirectly by the designer's decisions through the large-scale change in the end-users' behavior. In the current literature, context variables are treated as uncontrollable variables. Designers' approach to uncontrollable variables is developing predictive models to predict these variables accurately during the product life cycle and design the product accordingly. In a vehicle design context, examples include predicting the time that customers drive in traffic versus highway in order to design the most efficient engine. In this work, we propose a framework for bringing some dynamic aspects of the usage context into the controllable variable space which paves the way for designing the product and the usage context simultaneously. This will not only help with mitigating the negative impact of a product on its usage context, but will also offer a tool to change the usage context in order to get the most out of the product. We build the framework upon the rich body of literature

on multi-disciplinary design optimization (MDO) followed by a case study of optimization of fuel efficiency and mass for design of a sedan car.

Second study focuses on endogenous factors that change the user's preference. We developed a framework using Reinforcement Learning (RL) to design products that are adaptive and able to change their attributes over time as they interact with the user. This helps with the design of a product that not only matches each user's preferences (mass personalization), but also changes its attributes to maximize user satisfaction during the entire product life cycle without the need for designer's input or any information about the cause of change. Using real data on preference change for the design of a variable stiffness prosthetic ankle (VSPA), we explored different design approaches including traditional methods along with the RL framework, and compared different KPIs such as life cycle cumulative preference for the design approaches. Results show the superiority of the proposed framework over traditional design methods.

In the last study, we follow the framework proposed in the second study and explore how different aspects of a reinforcement algorithm exploration/exploitation behavior impact user willingness for providing feedback to an adaptive engineered system. A pilot study of 29 participants was conducted using an adaptive office chair. Statistical analysis of the results shows that the desirability of the system impacts the reported user willingness to interact over long periods of time. However, experiment data did not support the hypothesis that responsiveness of the system makes a significant difference in user willingness compared to desirable unresponsive system states.

This dissertation will open avenues for exploiting the abundance of data and data gathering tools to better design products that not only match individual user preferences, but also react to the changes in user preference over product life cycle. Starting from the usage context, we propose a framework for designing the product and the usage context simultaneously. We then try to remove the designer from the loop of the product design by proposing a framework that lets the product cycle through the design stages starting from preference elicitation and ending in the final product without designer's supervision. This helps with designing adaptive products that respond to user preference and change accordingly. The dissertation ends with exploring the factors affecting user willingness to repeatedly interact with the adaptive system.

CHAPTER 1

Motivation and Overview

1.1 Introduction

Driven by consumer demand and enabled by ubiquitous computation and connectivity, smart engineered systems are changing the way we live, work, learn, and play [1]. In this increasingly data-driven world, companies compete on new sources of data as well as new methods to gain insight into the product-user-environment system in order to develop products which offer superior user experiences. Many technology companies across diverse sectors such as advertising (Google), e-commerce (Amazon), and transportation (Uber), tailor services to individual users by leveraging data from past interactions to continuously update their products. For example, as they collect more data about how user preferences are changing over time, Amazon will change how its website looks, what products are on the front page, and even how much they cost [2]. The mass personalization and customization present in almost all software and Internet services results in high levels of continued user engagement and use [3].

The ability of software to react to fast-changing user preferences is in stark contrast to the state of physical products. Research has long shown that a significant percentage of consumer products are purchased and then never used [4]. From the unused treadmill to the specialized kitchen appliance for making poached eggs, some reports have found up to $10,000 dollars worth of unused products in the average American home [5]. This may be due in part to how physical consumer products are designed and produced. Current engineering design methods often focus on maximizing likelihood of purchase and focus less on what happens afterwards. As more smart products, such as subscription-based services like Peloton or HP printers, capable of personalization shift to business models in which revenue is driven by usage, designers will have to improve their ability to design for the entire product lifetime. As in software design, smart physical systems which adapt to dynamic user preferences can offer a competitive advantage to the organization and increased value to the end user. This dissertation investigates strategies for designing adaptive hardware systems which utilize user-generated data to personalize the product as well as react to user preference changes across the product life cycle. By enabling design teams to create improved adaptive hardware, this work seeks to increase user satisfaction over a product's lifetime and potentially reduce waste.

In traditional engineering organizations, a marketing team would elicit user needs and preferences, then a design team would create engineering specifications for a manufacturing and distribution team to use in producing and selling the new product [6]. Design research has offered a number of formal strategies for integrating these previously independent functions into a cohesive decision-making process [7, 8]. For example, Decision-based Design formulates organizational objectives and engineering constraints into a solvable optimization problem [9, 10]. These approaches have included marketing demand and cost models alongside existing engineering models. This type of analysis not only provides a systematic approach for managing trade-offs between competing objectives but also may mitigate miscommunication and other errors introduced at the intersection points between disciplines in traditional engineering organizations. Although these strategies have improved design outcomes, there are limited examples of design methods which can address situations with fast-changing dynamic user preferences. Existing design strategies are predicated on a paradigm consisting of an artifact with fixed functionality.

This dissertation proposes a future approach which relaxes this assumption. The work seeks to incorporate concepts from machine learning and software design to create smart adaptive hardware systems which leverage user feedback to automatically tailor their functionality to dynamic user preferences. The studies formalize methods for both eliciting user preferences over the product life-cycle and incorporating this new flow of information into the design process on a continuously updating manner.

1.2 Contribution and Chapter Overviews

This dissertation draws from the bodies of work in engineering design research. This literature includes research on formal design methods used to design products based on users' needs, preferences, and usage-context. Furthermore, this study builds upon the rich body of literature on machine learning methods and data-driven decision making. This work identifies and addresses the gaps in engineering design methods and proposes design frameworks for creating products that not only adapt to individual level preferences but react to the changes in users preference as well as the usage-context.

This dissertation begins by providing background on the use of user preferences in engineering design practice. It then gives a brief overview on design practices prevalent in other disciplines communities and identifies the gaps in more traditional engineering design methods that can be filled by bringing in the methods from software development (Chapter2). In the first study, the focus is on creating design frameworks for situations where the product usage-context is impacted by large-scale product adoption. A framework based on Decision-Based Design (DBD) and optimization methods is then proposed to address this type of design problems. A vehicle design problem was presented as a case study to showcase the value of the proposed framework (Chapter 3). Using a variable stiffness prosthetic ankle (VSPA), the second study investigated the application of reinforcement learning algorithms in design for dynamic user preferences, a topic that is not well explored by traditional engineering design methods (Chapter 4). The final study focuses on user willingness in providing prolonged feedback to adaptive devices. Successful implementation of reinforcement learning algorithms in adaptive design requires constant flow of information from the user in the form of preference feedback. This study sought to identify the factors and quantify their impact on users' willingness for providing feedback (Chapter 5). Together these studies propose new frameworks for designing products that react to user preference changes as well as optimizing the system level impact of the product leading to better user satisfaction throughout the product life-cycle. The dissertation is concluded with a discussion on how these findings add to the current state of work along with the limitations and the implications of the results (Chapter 6).

CHAPTER 2

Background

2.1 User Preferences in Design

Engineering design tasks require teams to translate abstract user needs into concrete engineering specifications [11]. As part of this process, it is critical for designers to address user needs and preferences in order to develop innovative and successful products which users value. Chen, et al., identify preference elicitation and incorporation of user preferences into formal design methods as key challenges for engineering teams [12]. Although user preferences are defined differently across disciplines, in this dissertation preferences will be defined as the rank ordering by desirability of alternative outcomes to a decision [13]. This choice is consistent with the prior work in design optimization this work is based on. By extension, user preferences will therefore be defined as the rank ordering by a user of the desirability of product alternatives [12]. Furthermore, dynamic user preferences are defined in this work as user preferences for the same set of product alternatives which vary over time. To motivate the proposed approach and situate this dissertation's contributions, the following sections will describe current qualitative and quantitative methods for eliciting and incorporating user preferences into the engineering design process. Techniques from other disciplines, such as marketing and machine learning, will also be presented. The following chapters will draw upon this body of work to propose novel strategies for improving system performance in situations where user preferences are dynamic.

2.1.1 Qualitative Approaches

Qualitative techniques, such as observations, interviews, focus groups, surveys, and ethnographic research, are among the most commonly used to elicit user requirements as they offer the designer a rich nuanced understanding of user preferences [11]. These strategies depend highly on the ability of the designer to process the in-depth information gathered and translate it into technical requirements. These strategies are limited in their scalability due to the high cost in time and resources for each additional participant. Thus, these techniques work best when the selected participants are representative of a target user group. Extensions of this type of work include lead-user theory and participatory co-design in which identified users are integrated into later stages of the design process [14, 15].

To integrate captured qualitative data into engineering design processes, researchers and practitioners have developed a number of approaches. User personas offer a framework for synthesizing and aggregating qualitative information gathered from many users into a single profile which is more easily used during the design process [16]. User journey mapping helps designers place user experiences into a timeline to better visualize product impacts [17]. Matrix representations are another approach which help designers synthesize technical and qualitative information. There is a rich body of literature summarized by Chan et. al on the use of Quality Function Deployment to improve product performance through linking product attributes and user preferences [18]. The House of Quality method has been widely used for similar purposes [19].

These "front-end" design methods rely on the design team to synthesize nuanced user information into engineering specifications. They are therefore susceptible to biases and misunderstandings inherent to human decision-making. Additionally, they require an upfront investment of resources to gather the raw qualitative data. They are best in situations in which the target user population is not well understood and are often used to identify new product opportunities. Another limitation is that differences in analytical techniques can make it difficult to manage trade-offs between user preferences and other organizational objectives which are quantitative in nature, such as cost and engineering physical constraints [12].

2.1.2 Quantitative Approaches

In order to ease the integration with existing quantitative engineering models, a number of quantitative approaches for eliciting and incorporating user preferences have been developed. For example, in Taguchi Robust Design and Design for Six-sigma, meeting customer satisfaction is the objective of design decisions [20, 21]. For quantitative approaches, choice modeling is a common method used to capture or elicit preferences. The creation of an explicit mathematical function linking product attributes to user preferences is an essential part of these methods. Design optimization or other mathematical techniques are then used to integrate the preference functions into the design process.

Modeling consumer choice behavior consists of a set of analytical techniques which seek to explain and predict selections made by target users from a set of product alternatives. These can then be extrapolated to a user population to estimate the demand, or quantity of products likely to be sold. Discrete Choice Analysis (DCA) and Conjoint Analysis (CA) are the major techniques used to capture user choice behavior. DCA relies on existing data sets of historical user choices to generate the probability an option is chosen based on available alternatives [22, 23]. CA uses survey responses of randomly generated ranked-choice questionnaires to generate a demand function based on product attributes [24, 25, 26]. Researchers have developed a number of modifications to DCA and CA to improve design outcomes including: using machine learning to improve the prediction accuracy [27], incorporating shape and aesthetic information in addition to technical attributes [28, 29], and incorporating the impact of sustainability preferences [30]. Recent work into Design Analytics (DA) has used big data and machine learning techniques to improve the predictive power of the preference models [31]. Through cyber-enabled products, sensory devices, and the collection of usage data, DA can automatically identify product attributes which impact user choice behavior. All of these strategies seek to create an single accurate functional mapping between potential combinations of product features and user preference or utility.

Design optimization is a major area of research providing systematic approaches for making design decisions. Formal methods under this umbrella, such as Analytical Target Cascading and Decision-based Design (DBD), are based on incorporating the preference function into the objective function of an optimization of the technical attributes of a product [32, 33]. In these strategies, the value of the designed artifact is maximized considering both the user and producer preferences [34]. Quantitative models of organizational objectives and physical constraints are also incorporated into the optimization formulation. Optimization algorithms can then be used to select a combination of product attributes or technical requirements which maximize the estimated user preferences. Design optimization has been used in a number of applications from aerospace engineering [35] to consumer product design [36] to the design of food systems [37].

There are a number of limitations to this type of approach. First, choice modeling approaches assume there is a latent preference function that can be modeled. Prior work has shown that users can be inconsistent in their preferences and may be constructing their preference functions when presented with the choice set [38]. Estimating the user preference function is difficult and there are a number of sources of error regardless of the selected technique. This error can translate into products which are not desirable to the target user. Of particular interest to this work, these approaches assume that the estimated preference function does not change between preference elicitation and the production of the artifact. Situations in which user preferences change rapidly are difficult to address using these methods.

2.2 Dynamic User Preferences

There are a number of areas in which user preferences for the same choice set vary with time. In fact, this idea drives the sheer variety of consumer product designs. As a user's preferences change, they will seek to purchase new and different products. For example, an individual may be faced with a choice between a sub-compact car and a larger sedan. The same individual may prefer and purchase the sub-compact car initially, and then several years later purchase the larger sedan. Although the sub-compact car was preferred to the sedan at the initial point in time, the individual's preferences changed over the time period. As this example illustrates, the definition of preferences is critically important to this dissertation as the following methods seek to accurately describe changes in the desirability rankings of alternatives, not explain why they changed.

2.2.1 Dynamic Preferences in Engineering Design

In response to applications where dynamic user preferences offer opportunities for improvement, designers and researchers have developed methods for addressing preference changes in specific cases. The simplest approach is to create multiple generations of the product. Iterative product cycles are central to traditional engineering practice. In this case, when target users change their preferences, the organization repeats the design process and produces a new offering [39]. Long development lead times have led to a variety of approaches which predict preference changes and design for future states. Design analytics extrapolates trendlines from current information to estimate future user preferences [40]. Uncertainty about future user preferences for technical performance can be mitigated through designing for system evolveability [41]. Similarly, modularity and flexible system design have been used to address potential changes in user preferences [42, 43, 44, 45]. Finally, in situations, such as ergonomic design, where user preferences are highly individual and hard to measure ahead of time, designers can give users the ability to customize features [46]. For example, office chairs frequently offer users control of a number of tuneable parameters. The sheer range of existing applications and approaches demonstrate how important addressing dynamic preferences is to successful product design.

2.2.2 Dynamic Preferences in Other Disciplines

Other disciplines have also addressed dynamic user preferences in creating adaptive systems. In marketing, organizations have been using Bayesian non-parametric approaches such as Gaussian Process Propensity Models, to estimate how individual users will engage with future product offerings based on past interactions [47]. These models can be used to help make ongoing decisions marketing resource allocations as they update as new information is obtained. Similarly, dynamic web and software design has long been used to change the content and appearance of websites and programs based on user interactions [48]. In robotics, intelligent control [49] and human-in-the-loop control [50] enable physical systems to respond to changing environments and human decisions. One commonality across all of these disciplines are systems which address dynamic user preferences automatically without the intervention of the designer. Dynamic websites, marketing dashboards, and intelligent control systems all adapt their functionality and performance based on changing preferences. These systems respond and adapt as the user continues to interacts with them, resulting in mass personalization and functionality which is responsive to user preference changes. This fundamental difference between these approaches and the design methods above is described in the machine learning literature using the concept of offline and online learning [51]. In machine learning, an offline learning system makes a decision based on an existing data set. For example, a classifier could be trained on a large set of observations and then used to make future decisions. By contrast, online learning systems make decisions as data arrives. In this case, a classifier is updated with each observation made and then used to make a decision before more information is collected. In computer science, algorithms exist on a spectrum between offline and online and make trade-offs between data storage constraints, data collection, and algorithm performance. With respect to user preferences, engineering design processes lie generally at the offline end of the spectrum, Fig. 2.1. User preference information is captured, incorporated into design decision-making, and an artifact is produced and sold to the user. By incorporating these techniques into design optimization, this dissertation hopes to move engineering design towards a more online process. Systems could gather user preference information and make more design decisions automatically without intervention from the designer.

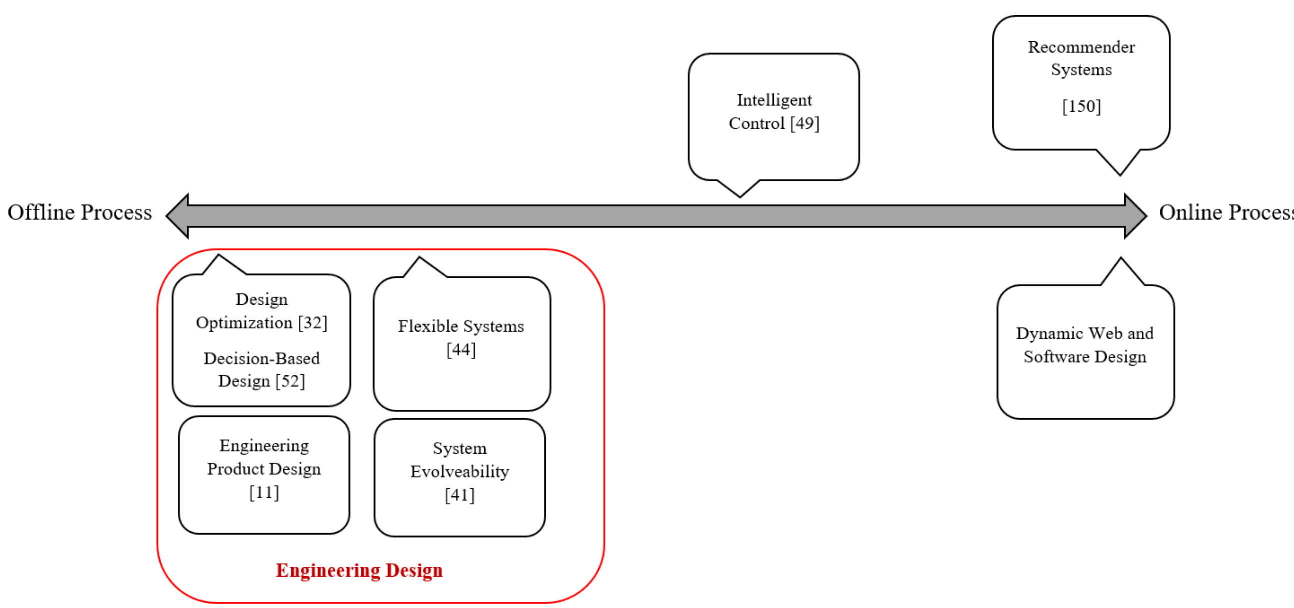

Figure 2.1: Incorporation of user preferences into product design.

2.3 Research Gap

Engineering design processes offer formal strategies for translating abstract user needs into concrete engineering specifications. As part of this process, engineers are required to elicit user preferences for product attributes and then subsequently integrate these preferences, along with organizational objectives and constraints, into their decision-making process. Dynamic user preferences pose a challenge to most existing methods. Work in other disciplines such as software design, controls, and machine learning offer new strategies for enabling systems to respond to changing user preferences. Machine learning identifies systems as either "offline," making decisions on gathered data, or "online," making decisions as data arrives. Current engineering design processes could be categorized as "offline" with respect to user preferences. This work seeks to bridge the gap between offline design optimization processes and the online systems used in other disciplines. In doing so, this dissertations lays the foundation for the development of systems which can automatically adapt to changes in user preferences due to either exogenous or endogenous sources. This systems could possibly offer greater user satisfaction over the lifetime of the product and enable mass personalization across a user population.

CHAPTER 3

Extending Usage Context-Based Design to Coupled Usage Contexts

This chapter was coauthored with Jesse Austin-Breneman and set to be published by ASME IDETC 2021 under "Extending Usage Context-Based Design to Coupled Usage Contexts: A Vehicle Design Case Study."

3.1 Abstract

Engineers must consider the usage context of a product in order to both predict its technical performance and model customer preferences. An emerging body of work in Decision-based Design (DBD) has elaborated various approaches for modeling the usage context in order to better predict customer choice behavior and select optimal product attributes. Building on this prior work, this study proposes a new method for formulating DBD problems in which product attribute values can change contextual factors. Results from a vehicle design case study demonstrate the utility of the proposed method for understanding how phenomena such as the rebound effect and induced travel demand connect system-level outcomes to design changes. This study suggests that the Design for Coupled Usage Contexts framework is a promising tool to further explore as a way to support designers making decisions which involve these types of mechanisms. Further exploration should include additional case studies to investigate other coupling mechanisms and design tasks.

3.2 Introduction

Successful new product development requires the effective integration of information from different disciplines across cross-functional teams in order to translate abstract user needs into concrete engineering specifications [6]. Model-based quantitative methods, such as Decision-Based Design (DBD), offer one approach for incorporating market research outcomes into technical engineering design [52]. For example, marketing techniques, such as discrete choice analysis (DCA) and conjoint analysis, can be used to construct a preference model of consumer choice behavior from the consumer survey data [34, 53, 27, 54]. These choice models can then be used to predict the future demand and market share of a product with a specific set of attributes. In DBD, enterprise-level objective functions based in part on these demand models drive design decision-making [8, 9, 10, 52]. Typically these strategies formulate design as an optimization problem [55, 56], enabling design teams to benefit from well-developed optimization techniques to quickly search a design space and find the optimal set of product attributes which maximizes demand or market share [32, 57].

In order to make accurate predictions of consumer behavior, the demand models used in DBD need to capture information about the operating context of the potential product. To understand whether the intended customers will choose a product over competing alternatives, DBD incorporates characteristics of the end-users and the market into the optimization problem setup. Customer profile attributes, such as gender, age, and income, are defined as the stable and permanent aspects of end-users which influence their choice behavior [58]. Building choice models based on the product attributes as well as the customer profiles result in models that better capture user preference heterogeneity which in turn yield more realistic demand models. Recent work on usage context-based design adds even more variables to the parameter space to better account for the preference heterogeneity caused by using products in different contexts. The impact of usage context on user's preference of a product is well explored and studied both by marketing community [59, 60, 61] and engineering design researchers [6, 62, 52, 58, 63, 64]. This work uses the He et al. definition of context variables as, "all aspects describing the context of product use that vary under different use conditions and affect product performance and/or consumer preferences for the product attributes [58]."

Despite the treatment of contextual factors as constant parameters, prior work has shown that designers can in some situations change user behavior in large-scale systems and possibly indirectly change the usage context and environment [65, 66, 67]. For example, sustainability researchers have identified the "rebound effect" in which increased demand for more energy-efficient technology negates the expected energy savings from improving the efficiency [68]. This phenomena is regarded by sustainability researchers as a negative consequence of good decisions and policies that were made without taking into account the macroscopic impact of the designers' choices on the socio-technical system of the end-users and their environment. In some scenarios, these phenomena may also negatively impact the utility of the product for the end-user. For example, increased use of more energy-efficient light bulbs may not translate to the user-desired savings on the electricity bill. This work proposes a framework for formulating DBD in cases like this, where the usage context and the product attributes are coupled and the usage context affects the consumer's utility.

The framework would enable optimal selection of product attributes by connecting the environmental states and user behavior to product attribute values. Designing in this manner may enable engineers to make more accurate predictions of user choice behavior by understanding interactions between the product and usage context. This changes the role of the designer from reactive to proactive with regards to the product usage context by bringing some aspects of environment and usage context from parameter space to the design variable space. This study presents a design framework for the optimization of this type of system and examines the conditions under which it would improve performance. The design of a passenger car is presented as an illustrative case example.

3.3 Related Work

This paper draws on design research literature as well as incorporating work from a number of disciplines such as economics, sustainability, and public policy on mechanisms by which decisions can affect the operating environment.

3.3.1 Decision-based Design Research

There is a rich body of literature modeling design as a decision-making process which fall broadly under the umbrella of Decision-Based Design (DBD) [8]. In DBD, product attributes, X, are selected by optimizing the expected utility, $E(U)$, of an enterprise decision-maker. Built on decision theory and economic models, this process connects marketing and engineering domains through cost and demand models [52]. The demand model consists of a discrete choice analysis model which compares potential attribute configurations to competitive alternatives and predicts the quantity of products sold for a given configuration. The generalized demand model incorporates market information drawn from surveys or focus groups on the demographics of target users, customer-desired attributes, engineering attributes, and market alternatives. By combining a demand and cost model, enterprise decision-makers can systematically select product attributes to maximize expected economic benefits or other potential objectives. DBD has been applied across a number of design domains including: aerospace [69], product-service systems [70], electric vehicle design [71], and supply chains [72]. It also serves as the basis for formal design methods such as Design for Market Systems [73, 74], and Design Analytics [40]. There are also a number of fundamental critiques of DBD, including the use of Utility Analysis and related assumptions [75]. This paper extends DBD research on contextual factors and therefore uses concepts and definitions from this literature throughout. Thus, this work assumes that design is a decision-making process occurring within an enterprise.

3.3.2 Context-Based Design

DBD is fundamentally a model of consumer choice behavior. The framework seeks to select product attributes that influence consumer preferences in a competitive market. One emerging research area focuses on how the usage context influences consumer choices. In these cases, consumer preferences for product attributes can vary widely depending on the environment in which the product will be used. This has long been studied within the broader design research community. Concepts such as "use environment" [6], "product design context" [62, 76], and "usage context" [77, 78] all describe the influence of the contextual factors on the desirability of engineered products. He et al. formalize usage context into a DBD process and include usage context into the demand model as an additional parameter set similar to user demographics [77]. This paper defines usage context as in He et al., categorizing usage context factors into the physical surroundings, social surroundings, temporal perspective, task definition, and antecedent states [77]. Further work by He et al. uses social network analysis to identify and measure social surroundings [63]. Wang et al. similarly used network analysis to examine interactions between a customers and products [64]. In all of these cases, usage context may differ across users or in time, but is independent of product attributes. For many design problem formulations, this treatment of usage context is valid and effort is better spent improving other aspects of the model. However, this paper seeks to address this issue for design tasks in which clearly identified mechanisms exist coupling usage context to product attributes.

3.3.3 Coupling Mechanisms

This paper draws upon several fields to provide examples typifying usage context dependence on product attributes. Since the illustrative case in this paper centers on automotive design, the following mechanisms are drawn from this area. However, these should simply be taken as examples of a larger class of scenarios. For example, the "rebound effect" is a widely studied phenomena in sustainability literature [79]. Efforts to improve energy efficiency to mitigate greenhouse gas emissions and gas consumption can lead to reductions in the cost of energy services and a corresponding increase in the use of energy. This effect can decrease or even erase any sustainability gains made by increasing efficiency [80]. Similarly, increased efficiency can lead to lower costs for vehicle travel, increasing the amount of travel for each vehicle [81]. This "induced travel demand" affects the overall usage context by impacting the traffic conditions and behavior of other vehicles. In this literature, demand refers to the total usage of all vehicles. Increased travel demand would reflect a user population's choice to drive their car more. The induced travel demand can reflect consumers' choices to switch modes of transportation for a trip, for example from public transport to private vehicles, or to make trips they might have otherwise not taken. By contrast, the demand model in DBD predicts the total number of cars of a particular type sold in a defined market. A related phenomena describes changes to driving behavior with the advent of driver assist technology. Researchers demonstrated drivers adapted their behavior after using different advanced driving assistance systems [82]. Observed drivers reduced the distance between themselves and other cars after using an adaptive cruise control system, even when not in a car with the assist system. This change would have major implications for traffic and safety conditions at larger scale. This paper uses these examples to create models linking product attributes such as fuel efficiency to usage context factors such as observed traffic and induced travel demand. By incorporating dynamic usage context factors, the DBD demand model may be able to make more accurate predictions of the number of products sold for a given attribute configuration.

3.3.4 Research Gap

Decision-based Design processes select product attributes to maximize an enterprise utility function. Researchers have adapted the general DBD framework to be suitable for a number of applications. This includes explicit representation of usage context factors to evaluate performance of design decisions under different operating conditions. However, as engineered systems grow increasingly interconnected there are a number of situations in which product attributes can affect usage context factors. Prior work across a range of disciplines has identified mechanisms by which product attributes and the operating environment are highly coupled. Because traditional DBD problem formulations view contextual factors as parameters which are independent of design decisions, they are limited in their ability to examine this type of scenario. This paper seeks to fill this gap by proposing a DBD framework in which the usage context factors are functions of the design variables. Given this background, the researchers seek to answer the following research questions:

1. How should designers formulate a DBD optimization problem in cases where the design attributes affect the usage context factors?

2. What is the potential impact on the optimal design of modeling these types of coupling mechanisms?

To answer these questions, this paper proposes a new framework, Design for Coupled Usage Contexts, and presents an illustrative case examines the impact of including these coupled contextual factors on the Pareto Solution Sets for previously studied problems in automotive design.

3.4 Proposed Framework

The proposed framework builds on the Usage Context-based Design (UCBD) framework by He et al. [78] in which the design process is formulated as a market share maximization problem. UCBD is an extension of Decision-Based Design that includes exogenous contextual variables in order to improve the predictive power of the utility and choice models. In this framework, a product is defined by a set of attributes and is sold to target user population for use in a given usage context, defined by a set of context variables. The objective of the enterprise is to maximize market share, which is modeled using discrete choice analysis (DCA) [83, 84]. Therefore, the maximization problem can be translated into maximizing the probability of a product being chosen by the users from a set of alternatives.

The solving process of the maximization problem starts with the system optimizer picking a set of values for the vector of design variables, $\mathbf{X_i}$. The analysis functions, $f(\mathbf{X}, \mathbf{E}, \mathbf{S})$, use the picked set of design variables along with context variables ($\mathbf{E}$) and customer profile attributes ($\mathbf{S}$) to generate the corresponding product attributes, $\mathbf{A_i}$s, Eq. 3.1. Product attributes can be described as the perceived characteristics of the product by the users [52]. Since the attributes are perceived by the users, they may depend on the context, and the customer attributes as well as the design variables. For example, the wheelbase of a vehicle is a design variable, whereas the handling of the vehicle can be described as its attribute perceived by a specific user in a specific context.

$$\mathbf{A_i} = f(\mathbf{X}, \mathbf{E}, \mathbf{S}) \tag{3.1}$$

After generating all the product attributes using design variables, context variables, and customer attributes, a utility model estimates the utility of the set of attributes given the usage context of the product, and the customer profile attributes, Eq. 3.2.

$$W_i = W(\beta : \mathbf{A_i}, \mathbf{E}, \mathbf{S}) \tag{3.2}$$

Regression based models are often used to build the utility function in DCA [52], where W_i is the model for observed utility of product i having a set of attributes. The model is based on assuming that users have a utility function consisting of an observed part W and a random error term ϵ. Equation 3.3 shows the utility of product i.

$$U_i = W_i + \epsilon \tag{3.3}$$

Using the calculated utility of the product, W_i, and the calculated utility of the alternative products, W_js, the system optimizer calculates the probability of the product i being chosen from the set of n product alternatives using a multinomial logistic regression model, Eq. 3.4. The value of the calculated probability for the current set of design variables is then compared with the previous probability values, and the next set of design variables is chosen accordingly. The entire process is repeated until some stopping criterion is reached.

$$\underset{W_i}{\text{maximize}} \quad P_i = \frac{e^{W_i}}{\sum_{j=1}^{n} e^{W_j}} \tag{3.4}$$

Due to multi-disciplinary nature of this optimization problem, it belongs to a large class of optimization problems called multidisciplinary design optimization (MDO) problems solved using MDO techniques [32]. MDO techniques work by decomposing the optimization problem into separate units and define the information flow between the units. These decomposition methods are called MDO architectures and a handful of them have been proposed in the literature. For UCBD framework, all-at-once (AAO), individual discipline feasible (IDF), and multidisciplinary feasible (MDF) architectures can be used to solve the optimization problem [85, 32]. Figure 3.1 summarizes the problem formulation.

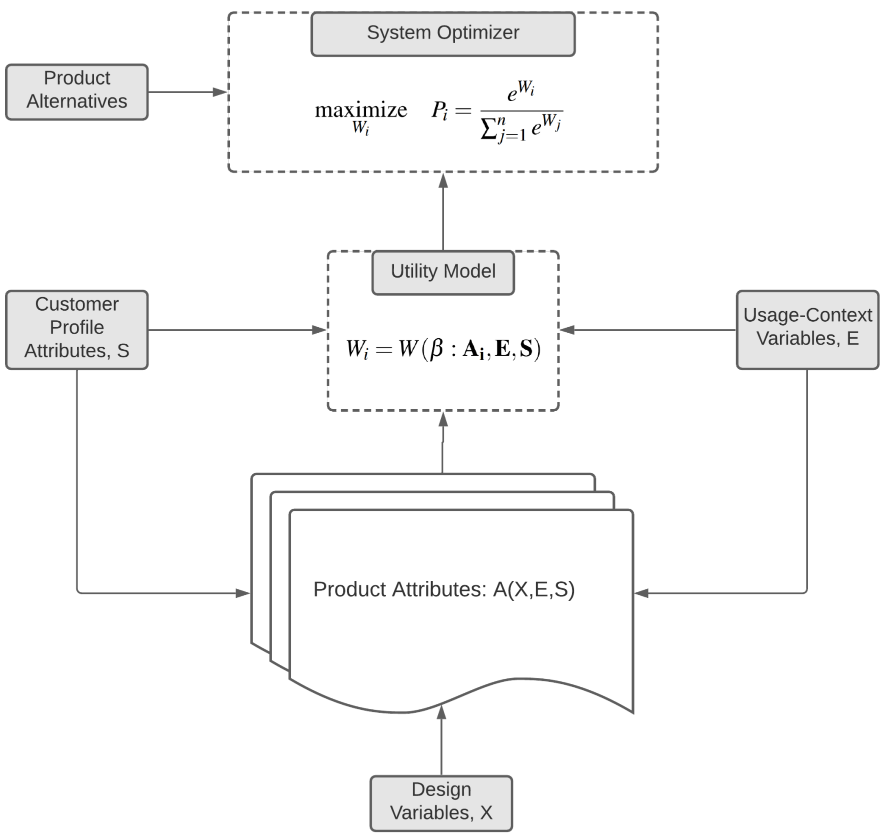

Figure 3.1: Design problem formulated as an optimization problem.

3.4.1 Evolution Function

In order to improve the prediction accuracy of the usage context variables in presence of coupling mechanisms, the impact of product attributes on usage context variables has to be modeled explicitly in the optimization problem setup. So far in the problem formulation, usage context variables, $\mathbf{E}$, are treated as exogenous parameters that are not impacted by the system consisting of the users, the product and the environment they interact with. To support decision-making in cases similar to those discussed in the previous section in which usage context is impacted by the product, an evolution function is introduced to the problem formulation to capture the possible system level effect of the product on the usage context. The evolution function is a function of product attributes $\mathbf{A}$, design variables $\mathbf{X}$, and customer profile attributes $\mathbf{S}$ (or target market characteristics) that outputs the set of new usage context variables for the evolved system of users and their environment once the product is introduced.

The new problem formulation replaces $\mathbf{E}$ in Eq. 3.2 and 3.1, with evolution function, $\mathbf{EV}(\mathbf{X}, \mathbf{A}, \mathbf{S})$, as shown in Eq. 3.5 and 3.6.

$$W_i = W(\beta : \mathbf{A_i}, \mathbf{EV}(\mathbf{X}, \mathbf{A}, \mathbf{S}), \mathbf{S}) \tag{3.5}$$

$$\mathbf{A_i} = f(\mathbf{X}, \mathbf{EV}(\mathbf{X}, \mathbf{A}, \mathbf{S}), \mathbf{S}) \tag{3.6}$$

Note that the evolution function is a function in general form that includes the scenarios where usage context variables and environment are not affected by the product design as formulated in the previous subsection. Figure 3.2 summarizes the new problem formulation.

The new optimization problem can be formulated in multidisciplinary feasible (MDF) architecture and solved using fixed point iteration [32]. In this problem setup, first a set of design variables that are feasible across all the disciplines (attributes) are found using some initial value for the context variables. Then using the evolution function and the newly found multidisciplinary feasible design variables, the new set of usage context variables are found. The new context variables are then again fed into the attribute functions along with the same set of design variables. The process is repeated until all the context variables and attributes converge (fixed point iteration). The converged variables and attributes are then sent to the system optimizer to be evaluated. The system optimizer evaluates the objective function and picks the next set of design variables to be explored. The steps are repeated by the system optimizer until some stopping criterion is reached.

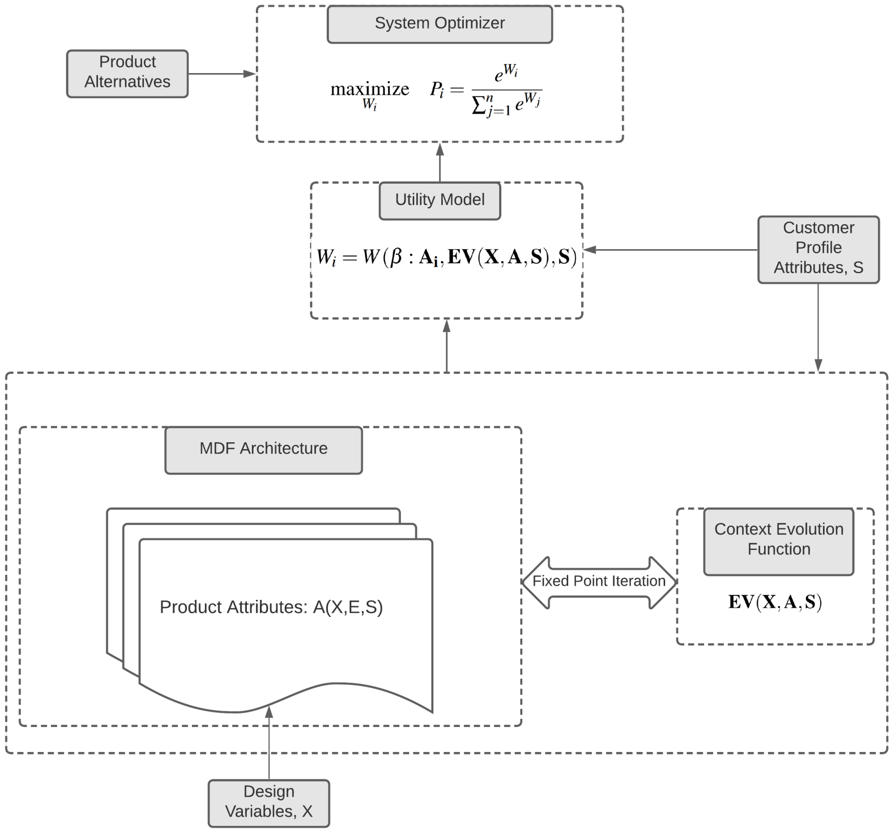

Figure 3.2: New design problem formulation with the evolution function.

3.5 Illustrative Case Study

The proposed framework is implemented in the context of the design of a sedan passenger car. Using a large-scale high-fidelity traffic simulation platform [1], the system level impact of the product through usage context variables on its optimal fuel economy and safety design is investigated.

3.5.1 Simulation Platform

To capture the complexity and emergent phenomena arising from having a complex large-scale system, an agent-based model [86, 87] of traffic with realistic agent behavior is developed in Python. A relatively large area in the southeast Michigan covering multiple cities is simulated. Figure 3.3 shows the simulation area which is fed to the platform from Open Street Map[2]. Figure 3.4 shows a snapshot of simulation for the area shown in Fig. 3.3.

To validate the simulation platform, first a realistic demand model is developed as the model

[1] Simulation Platform details are explained in Appendix A
[2] https://www.openstreetmap.org

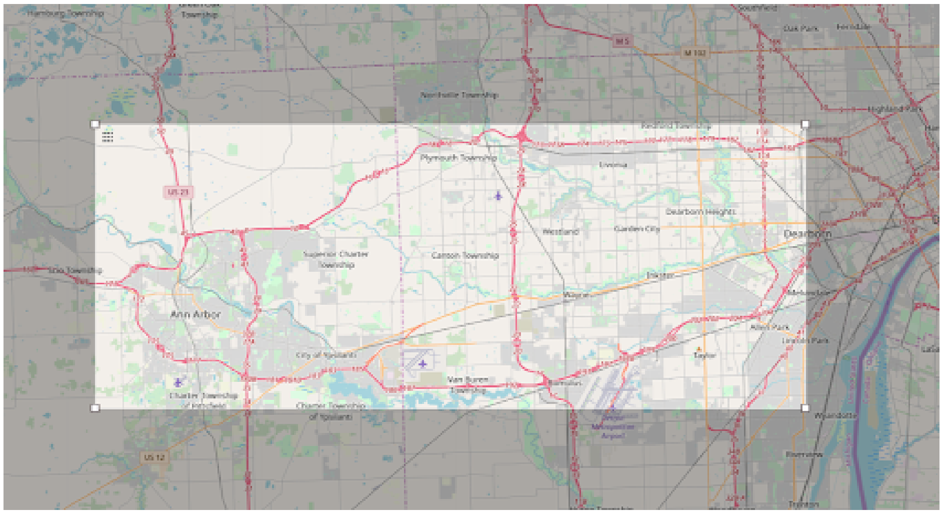

Figure 3.3: Simulated area in southeast Michigan.

input based on a combination of datasets from 2010 U.S. Census and the Southeast Michigan Council of Governments (SEMCOG) [3]. The outputs of the simulations were then compared with the available real world data and transportation metrics. The results are close to the transportation metrics reported for the United States and SE Michigan. For example, as the main purpose of the model to predict fuel economy, the simulation predicts an average mpg of 23.5 which is within 6% of the average mpg of 24.9 reported by EPA [4].

3.5.2 Problem Setup

Following the setup of DBD framework, the problem is formulated as a maximization problem based on a multinomial logit model (MNL) estimation, Eq.3.4. The goal of the optimization problem is then to find the optimal combination of attributes that maximizes the probability of a product with certain set of attributes being chosen by the end-users. Furthermore, the problem can be reduced to the maximization of the linear predictor (observed utility, W) of the MNL model Eq. 3.7.

$$W(\mathbf{A}) = \beta_1 A_1 + \beta_2 A_2 + ... + \beta_m A_m \tag{3.7}$$

Where A_i and β_i are attributes and MNL coefficients, respectively. The linear form of the equation makes the optimization process relatively simple. The best product is the one that have the largest possible A_is corresponding to positive β_is, and smallest A_is corresponding to negative β_is regardless of the size of the MNL coefficients. At first glance, this may render the process of

[3]https://semcog.org/

[4]https://www.epa.gov/automotive-trends/highlights-automotive-trends-report

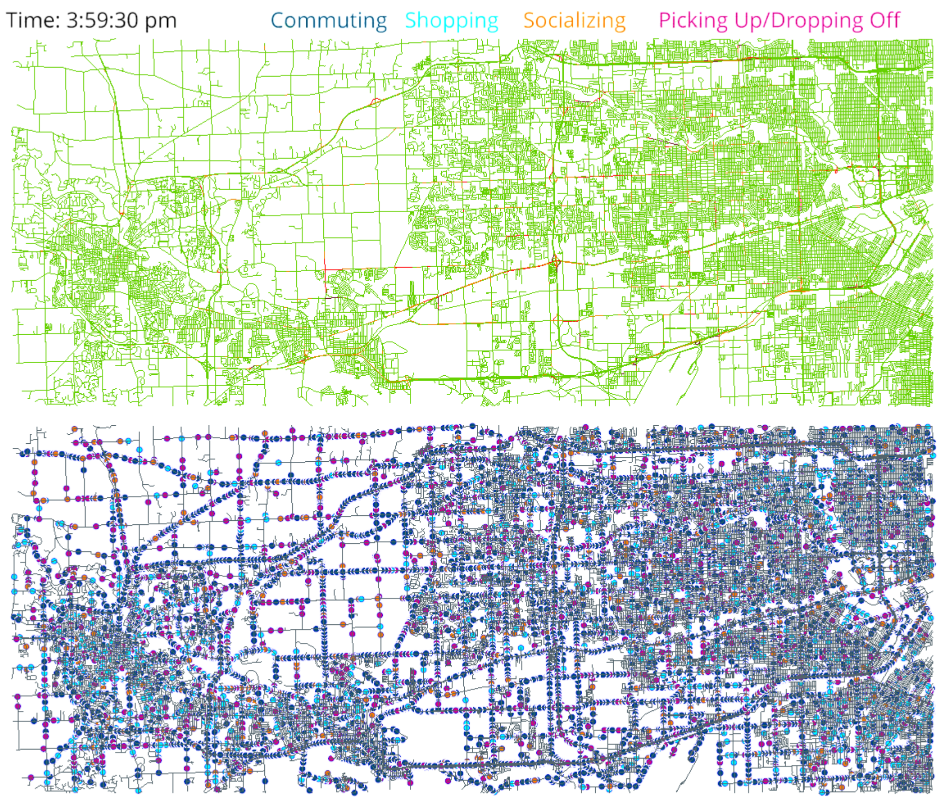

Figure 3.4: Simulation snapshot.

accurately estimating MNL coefficients less significant, for as long as the sign of coefficients are accurate, the optimal product can be easily found. However, in many cases the attributes conflict with each other. For example, designing a high performance engine which is mostly desirable (positive β_i) adversely impacts fuel economy of the vehicle (also positive β_i). This will result in a trade-off between fuel economy and engine power. Although both attributes have positive coefficients, the optimal values for the two attributes depend on the size of the MNL coefficients.

This case study focuses on a different well explored trade-off, namely fuel economy vs. safety. The trade-off between vehicle safety, mass, and fuel economy is well studied in the literature [88, 89]. Heavier cars are on average safer but perform worse in terms of fuel economy. As the two competing attributes of safety and fuel economy have positive coefficients in the MNL model (both are desired by consumers), finding the best product design for a market requires estimating β_{FE} and β_{Safety} accurately for the target users, Eq.3.8.

$$W(MPG, Safety) = \beta_{FE} * MPG + \beta_{Safety} * Safety \tag{3.8}$$

Equation 3.8 can be further simplified by dividing the equation by β_{FE} and defining $\alpha = \frac{\beta_{Safety}}{\beta_{FE}}$. α depends on the target end-users and shows how customers weigh safety versus fuel economy.

$$\frac{W}{\beta_{FE}} = W'(MPG, Safety) = MPG + \alpha * Safety \tag{3.9}$$

The case study analysis is conducted under the assumption that the designed product will have a significant market share, and the impact of the products from competitors on the system is negligible. Figure 3.5 summarizes the problem formulation.

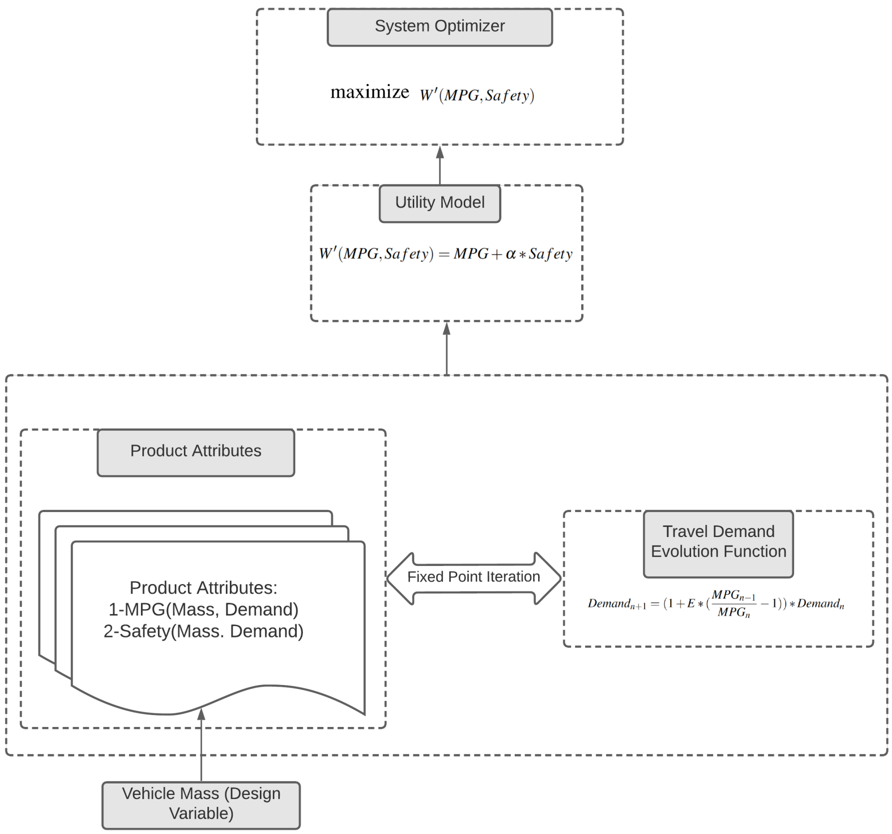

Figure 3.5: Case study formulated as an optimization problem.

3.5.3 Evolution Function

Having travel demand (total miles traveled not product demand) as the usage context for the evolution function, the evolution of travel demand is modeled using the well studied fuel cost elasticity, E, [68, 80, 81]. Equation 3.10 shows the relation between trip demand change and change in fuel cost.

$$Demand_{n+1} = (1 + E * \frac{FC_n - FC_{n-1}}{FC_{n-1}}) * Demand_n \tag{3.10}$$

Where E is fuel cost elasticity, n and FC are time step and fuel cost, respectively. Demand is a single number that works as a multiplier increasing or decreasing the number of trips generated by the demand model without changing the distributions used for generating the validation demand.

Under similar demand and travel conditions, fuel cost can be written as $FC = \frac{K}{MPG}$, where K is some constant. Substituting in the Eq. 3.10:

$$Demand_{n+1} = (1 + E * \frac{\frac{K}{MPG_n} - \frac{K}{MPG_{n-1}}}{\frac{K}{MPG_{n-1}}}) * Demand_n \tag{3.11}$$

Canceling constants K, the evolution function in terms of product attribute MPG is found, Eq. 3.12.

$$Demand_{n+1} = (1 + E * (\frac{MPG_{n-1}}{MPG_n} - 1)) * Demand_n \tag{3.12}$$

This work uses the value -0.28 for E, estimated by Greene [80] based on U.S. national time series data on vehicle travel by passenger cars and light trucks.

3.6 Results

In this case study, vehicle mass is the only design variable and fuel economy is found through simulation of the transportation system using vehicles of that mass. The following figures present the Pareto frontier [32] for fuel economy and safety for a range of α in Eq.3.9. Determining a single optimal design would require estimating α for a specific market and this is out of the scope of this paper. Corresponding optimal mass values are found for each α in the original and evolved systems. A comparison of the optimal mass values is plotted in Fig.3.10.

3.6.1 Creating Initial Pareto Frontier

Having mass as the only design variable, to form the Pareto frontier, two sets of data points are needed. The first set relates optimal mass of the vehicle to its safety. Under the assumption that automotive designers optimize the vehicle mass, real world data on average vehicle safety vs. mass is a good representative of optimal mass vs. safety curve for an average car. The data from Insurance Institute for Highway Safety (IIHS) on vehicle curb weight versus fatality rate for 2007-10 model cars is used [90]. Fatality rate is the number of deaths per million registration years, and vehicle safety is defined as negative fatality rate.

For the second set of data points, the same logic is followed and it is assumed that automotive designers optimize all the variables impacting fuel economy in the vehicle design process including power train efficiency, coefficient of drag, and idling fuel consumption. Using the average values of the mentioned parameters for a sedan car and the vehicle mass, the simulation calculates the average fuel economy in mpg. Running the simulation for a range of vehicle mass and using the (optimal) parameters impacting fuel economy, the second set of data points relating optimal mass to optimal fuel economy is derived. Finally, the two data sets are combined using mass as the intermediate variable to form the Pareto frontier of optimal fuel economy vs. safety (negative fatality rate). Figure 3.6 shows the Pareto frontier derived from real data and simulation for a sedan passenger car.

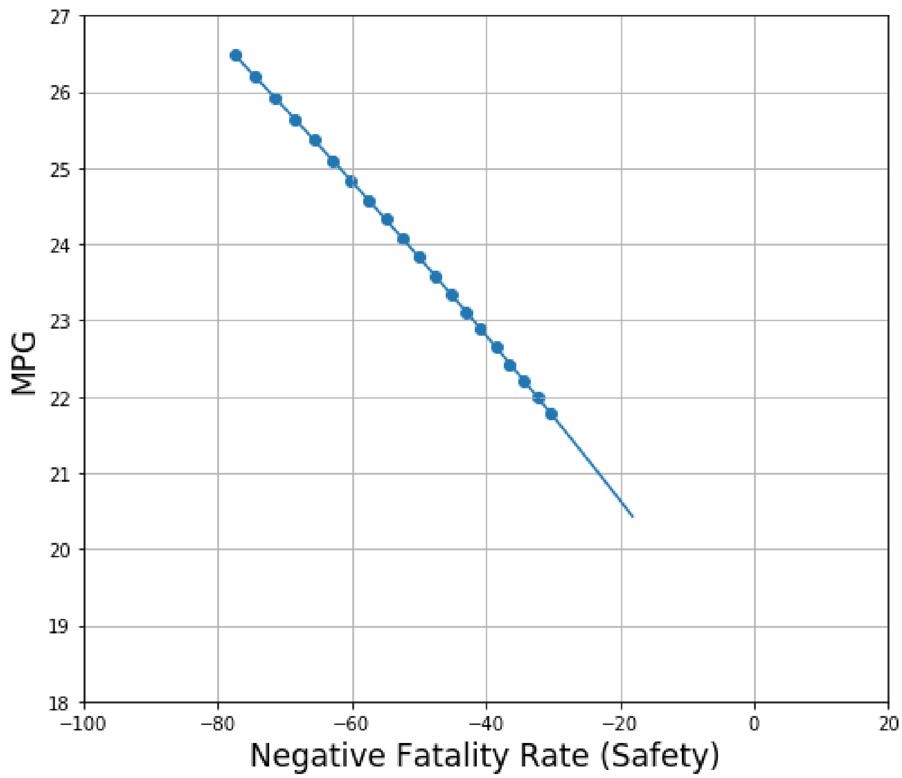

Figure 3.6: Pareto frontier of fuel economy vs. safety.

Optimizing Eq.3.9 for different values of α gives a different point on the Pareto frontier indicating the optimal trad-off between fuel economy and safety for a specific α value. The dots on the Pareto frontier in Fig.3.6 show the optimal trad-offs for a set of α values.

3.6.2 Evolving the System

Using the evolution function and the results from simulation platform, the future demands for all the design points (in this case vehicle mass) are predicted. The future demands are then again fed into the model to predict the next iteration's demand. The process is repeated until the demand for each design point (mass) converges to a number. Figures 3.7 and 3.8 show how demand and MPG evolve for a set of vehicle weights. Only number of trips are shown for the demand as the distributions used for generating trip types, times, starting points, and end points remain unchanged.

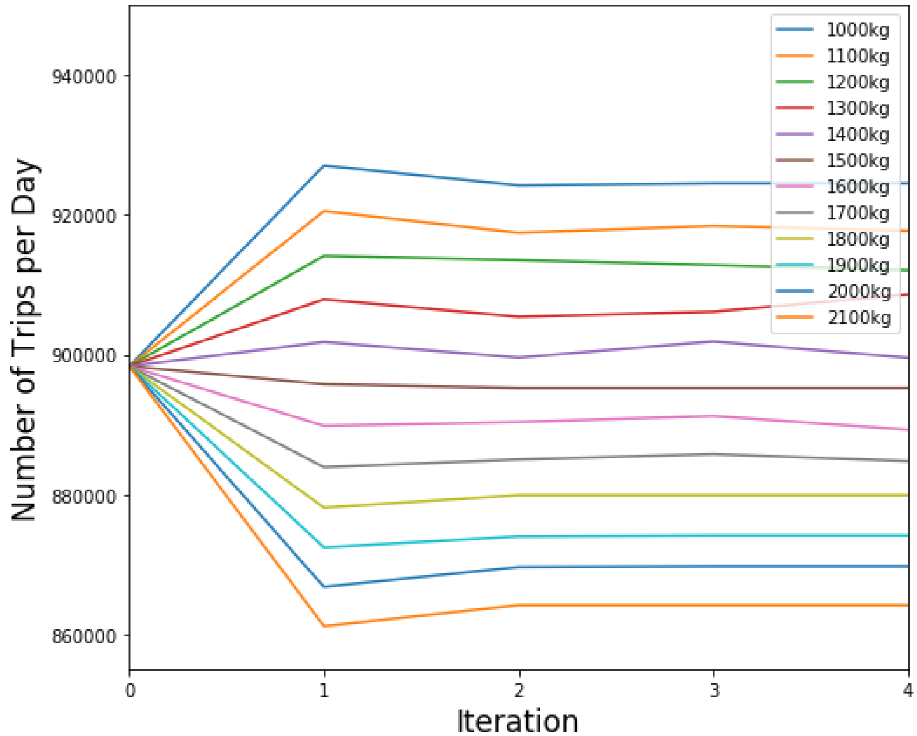

Figure 3.7: Evolution of demand for different designs.

Using the final values of MPG, Fig.3.9 shows the new Pareto frontier after the system passes the transition time along with the original Pareto frontier.

Similar to Fig.3.6, Fig.3.9 shows 20 dots on both original and final Pareto frontiers that correspond to the same 20 different values of α. The fuel economy/safety trad-off points for each value of α have changed for the final Pareto frontier compared to original Pareto curve. Since the optimal product design is of interest, namely optimal mass for each α value, optimal mass for the two Pareto curves are plotted in Fig.3.10 for a set of α values. The graph shows that the optimal mass for the evolved system is on average around 500 kg more for the same α compared to the original system.

3.6.3 Single Vehicle Crash Death Rate

This case study assumes that all vehicles in the transportation system are identical. Thus, assumption the designed passenger car has the same weight for all consumers. As studied by Evans et al. [88] and Tolouie et al. [89], one of the major causes of fatality in car accidents is the change of speed before and after an accident due to conservation of momentum. Conservation of momentum

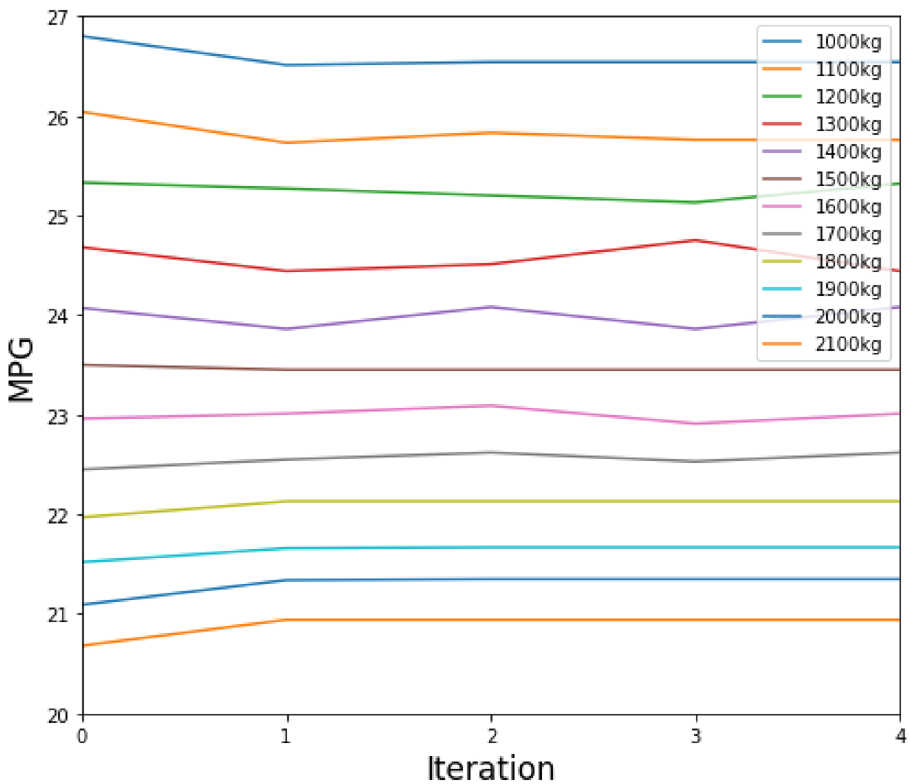

Figure 3.8: Evolution of mpg for different designs.

causes a larger speed change hence higher chance of fatality in the lighter car in an accident. The data from Insurance Institute for Highway Safety (IIHS) used above consists of all types of accidents, namely multi-vehicle crash between vehicles with different weights, single vehicle crash, and roll over. However, assuming identical vehicle mass causes an overestimate of fatality rates for lighter cars and an underestimate for heavier cars in this analysis. Therefore, additional analysis of the fatality rate for a given mass was performed using the subset of IIHS data involving single vehicle crashes in which the difference in mass is not relevant. The corresponding optimal mass values for the original and evolved system for this safety metric are shown in Fig.3.11.

3.7 Discussion

Results from the vehicle case study demonstrate how the use of an evolution function results in different optimal designs for the same user utility function across a range of values of α. The case study shows that when total miles travelled and vehicle fuel economy are interdependent, the utility of the optimal design from original formulation can be improved regardless of the relative weighting of safety and fuel economy. In this scenario, optimizing the vehicle mass without con-

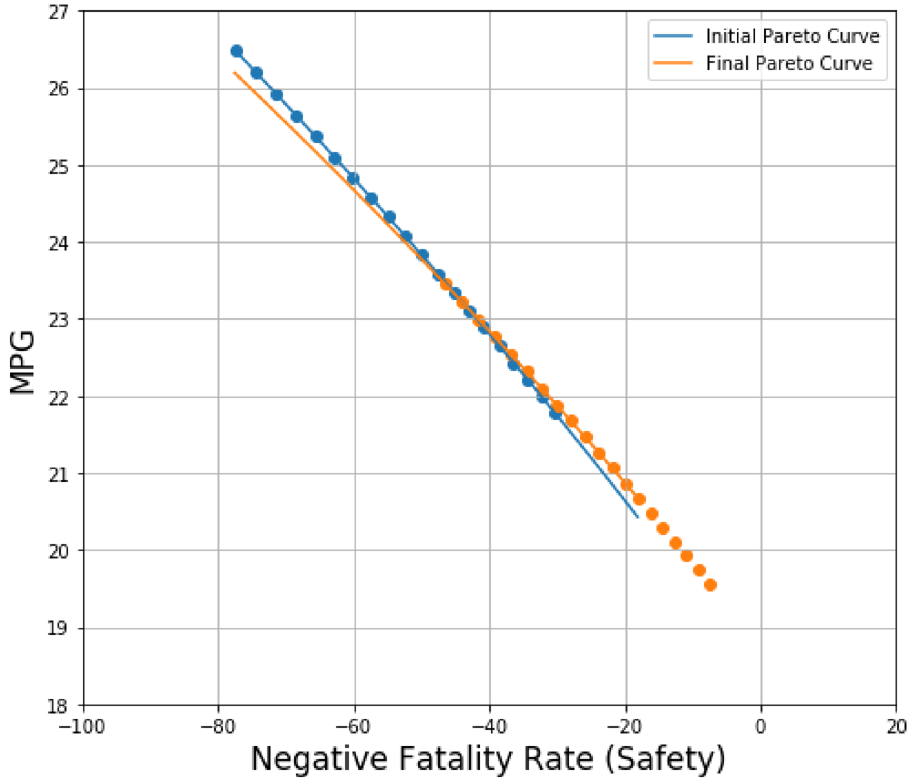

Figure 3.9: Original and final Pareto frontiers.

sidering the coupling between mass and travel demand will result in a vehicle which is lighter than necessary and therefore less safe. As seen in Fig.3.9, although the Pareto frontiers have similar shapes the performance of the optimal design for a given α changes significantly. Generalizing to other design tasks, these results suggest that incorporating an evolution function into a DBD framework can provide additional insight into system behavior and the impact of usage context on user utility.

The case study results correlate with expected system behavior. Starting from the initial Pareto frontier, the trade-off between fuel economy and safety is well captured in Fig.3.6, as with improvement in MPG, safety is negatively impacted. The same relationship is also observed in the evolved Pareto frontier, Fig.3.9. However, the slope of the new Pareto curve is slightly shallower compared to the original curve. The change in the shape of Pareto curve moves the trade-off points for each α value toward safer, and heavier designs. The optimal mass found from the two scenarios differ approximately by 500 kg on average for the same α as shown in Figure 3.10.

Although the magnitude of the change may be inaccurate, the overall trend in the difference between the initial and evolved Pareto curve can be explained by the evolution of travel demand (total miles travelled) Fig.3.7, and fuel economy, Fig.3.8. In the both formulations, lighter designs perform better in terms of fuel economy. However, in the evolved formulation, improved fuel

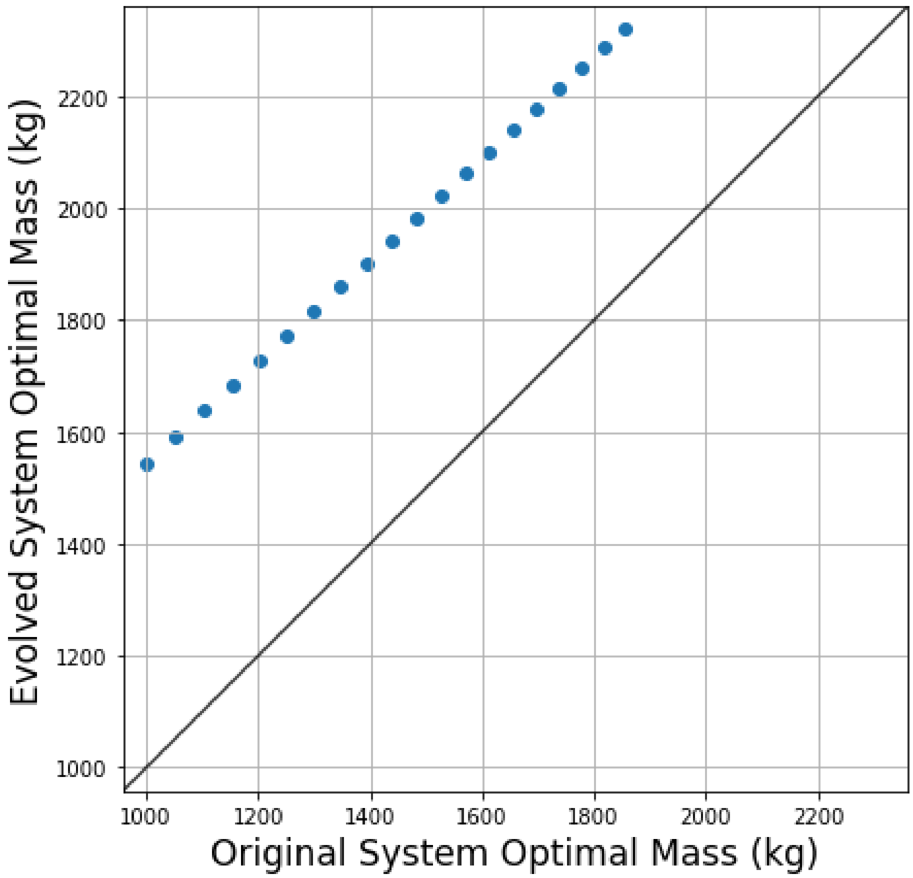

Figure 3.10: Optimal mass for original and evolved system for the same set of α values.

economy induces travel demand due to fuel cost elasticity. Induced demand increases traffic and cancels some of the fuel economy gains of lighter vehicles. For lighter vehicles, the two opposing forces of induced demand and reduced mass result in a fuel economy which is slightly worse than than what is predicted without taking into account the coupling between the product and the environment. For heavier vehicles, the opposite effect is true with heavier designs resulting in worse fuel economy, reducing demand and consequently traffic. The reduced traffic increases fuel economy and mitigates some MPG loss of the heavier car. The effect is greater at the extremes of the fuel efficiency range. This is seen in Fig.3.9, as the left tail of the new Pareto curve has the largest deviation from the original curve and differences become less pronounced towards the middle points.

Finally, as discussed in section 5.3, using real world data without considering the weight differences in an accident leads to fatality rate overestimate for lighter vehicles and underestimate for heavier vehicles. As shown in Fig. 3.11, although less pronounced, these results exhibit a similar system behavior with the original and evolved functions resulting in different optimal mass values.

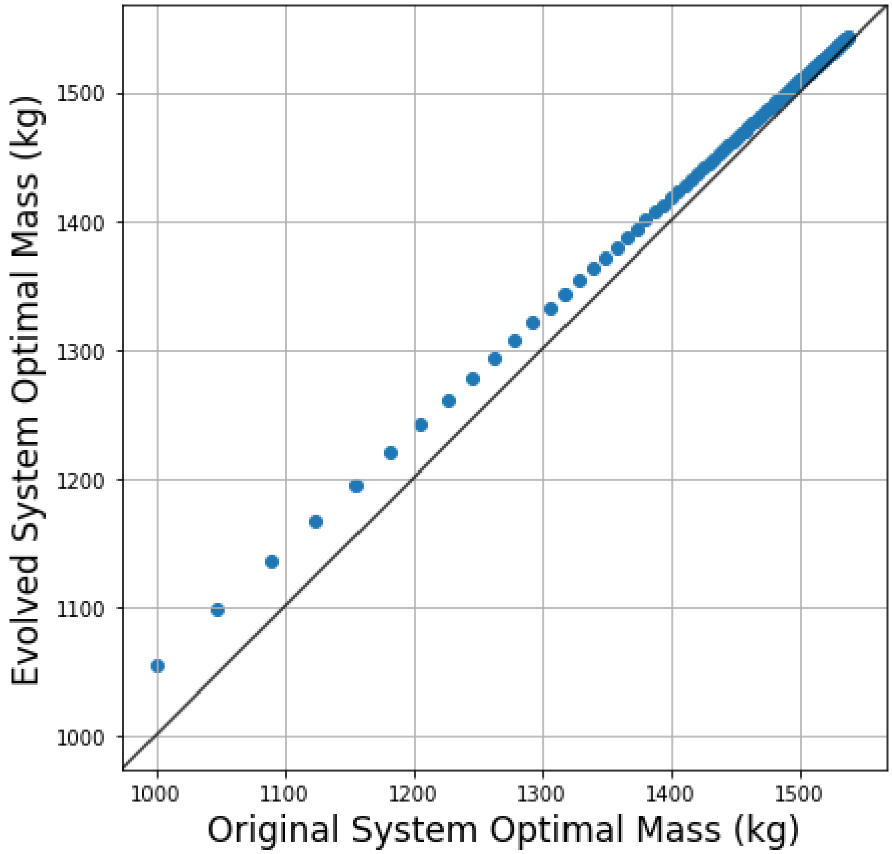

Figure 3.11: Optimal mass for original and evolved system for single vehicle crash death rate.

3.8 Conclusion

An emerging body of work in Decision-based Design has elaborated various approaches for modeling the usage context of a product in order to better predict customer choice behavior and select optimal product attributes. For example, He et al. used social network models to estimate how different market segments made purchasing decisions [63]. This work builds on these efforts by examining design scenarios in which product attributes can change contextual factors. The optimal set of product attributes therefore depends on the coupled interactions between the product and its context. The proposed framework and results from a vehicle design case study provided the following answers to the research questions.

1. How should designers formulate a DBD optimization problem in cases where the design attributes affect the usage context factors?

   The proposed framework incorporates an evolution function and iterative results from a simulation platform to examine how demand changes across various product attributes. This allows designers to explore the design space and investigate tradeoffs across emergent usage contexts.

2. What is the potential impact on the optimal design of modeling these types of coupling mechanisms?

   Results from the vehicle design case study suggest that the optimal design can change significantly from a traditional DBD approach to the evolved case. The evolved Pareto curve suggests that if you consider traffic due to induced demand, the optimal design will have a higher mass for the same relative weighting of mass versus fuel efficiency. The magnitude of this effect is not linear across the range of the design space and grows larger as the design point deviates from the mean.

In conclusion, the proposed method will enable engineers to understand trade-offs and find optimal designs while considering changes in operating conditions due to design changes.

3.8.1 Future Work

The results of the study are limited by two simplifying assumptions. First, in real world scenarios, designers are designing a product as well as competing with alternative products whose market shares and system level impacts are not negligible. Future work should therefore extend the problem formulation to include market competition by incorporating the framework in a game theoretic set-up where the evolution function is a function of the attributes of all the products with significant market share.

Finally, to show the framework's utility for real world design problems, this work used available data on automotive design as much as possible. This meant simplifying details on the exact relationships between different product attributes, in particular the impact of coupling of fuel economy and fatality rate on the evolution of the Pareto frontier. The degree of coupling between product attributes is high for many systems and is not dealt with in this case study. Nonetheless, the proposed framework is capable of handling any attribute coupling through the MDF architecture. For future work, a design case study with all the attribute coupling explicitly formulated should be explored using this framework.

CHAPTER 4

Theoretical Framework for Design for Dynamic User Preferences

This chapter was coauthored with Elliott Rouse, and Jesse Austin-Breneman and the results are published in [54].

4.1 Abstract

A key assumption of new product development is that user requirements and related preferences do not vary on time scales of the process length. However, prior work has identified cases in which user preferences for product attributes can vary with time. This study proposes a method, Design for Dynamic User Preferences, which adapts reinforcement learning (RL) algorithms for designing physical systems whose functionality changes with user feedback. An illustrative case comprised of the design of a variable stiffness prosthetic ankle is presented to evaluate the potential usefulness of the framework. Lifetime user satisfaction for static and dynamic design strategies are compared over simulated user preferences under a number of conditions. Results suggest RL-based strategies outperform static strategies for cases with dynamic user preferences despite significantly less initial information. Within RL methods, upper-confidence bound policies led to higher user satisfaction on average. This study suggests that further investigation into RL-based design strategies is warranted for situations with possibly dynamic preferences.

4.2 Introduction

New product development requires engineering designers to perform requirements engineering tasks to translate user needs into engineering specifications during the early stages of the design process. This process can be broadly described as occurring in three sequential steps: 1) elicitation of user requirements, 2) requirements analysis, and 3) requirement specification [91]. In this way the designer progresses from the abstract user domain to the concrete functional domain. A key assumption of this process is that user requirements and related preferences are independent of time on the scale of the length of the development process. Since the selection of product attributes is a critical decision in this process, this study uses the definition of user preference as the rank ordering by desirability to the user of possible combinations of product attributes [34, 92]. In using existing strategies, designers assume there is an underlying "latent" preference function for product attributes which will not change between when preferences are elicited and when the final artifact is produced and used.

Dynamic web design has partially relaxed this assumption in the software world. Instead of driving design decisions using previously elicited user requirements and preferences, dynamic software collects feedback throughout a user's interactions with the software. This feedback is used to update the functionality of the software as additional information is learned. The ubiquitous "like" button is a powerful tool for enabling a software product to tailor itself to an initially unknown "latent" preference function using Bayesian updating techniques. Depending on the algorithm selected, these updating techniques can also account for preference functions with dynamic behavior.

Although responsive systems and active control techniques, such as intelligent control [93] and human-in-the-loop control [94], allow for physical systems which adapt their functionality in response to time-varying feedback mechanisms, these are commonly based on objective measures and not subjective user preferences. This work proposes a framework for the design of physical systems which change their functionality in response to changing user preferences over time. *Design for Dynamic User Preferences* adapts the dynamic software paradigm and related techniques to the design of physical hardware.

The Design for Dynamic User Preferences framework could improve overall user satisfaction and reduce development cost under certain conditions. Current design methods use observation, interviews, or surveys to estimate a preference function for a user population to guide design decisions. In cases where the user population has sufficiently homogeneous time-independent preferences, these methods efficiently design for the target audience. Additionally, the unbiased post-purchase user feedback required by the proposed strategy can be difficult and expensive for the designer to obtain with current tools. However, there are many situations, such as the design of prosthetic limbs examined here, in which user preferences within a target population are highly heterogeneous, change quickly, and where it may be prohibitively expensive or practically infeasible to conduct a broad survey. In these cases, the proposed framework would enable both personalization to each individual and a dynamic response to changes in an individual's preferences over time more efficiently than current methods. By using feedback to tailor the design to an individual, the framework could improve user satisfaction over the lifetime of the product and potentially reduce development cost by requiring less information upfront. This study presents a design framework for the optimization of this type of system and examines the conditions under which it would improve performance. The design of a variable stiffness prosthetic ankle is used as an illustrative case example.

4.3 Related Work

This study draws upon literature in a variety of fields including engineering design research, controls, and machine learning.

4.3.1 Early Stage Design Processes

New product development tasks require designers to translate user requirements from the customer domain into engineering specifications in the functional domain for further design work [95]. Formal methods for eliciting and understanding user requirements fall under the broad umbrella of "early stage design" [96]. A key part of this process is the designer connecting the technical features and attributes of a product to user needs or requirements based on collected information. Methods to achieve this goal can be broadly categorized as falling into two areas: qualitative methods, and matrix representations, such as Quality Function Deployment and House of Quality.

4.3.1.1 Qualitative Methods

Qualitative techniques, such as observations, interviews, focus groups, surveys, and ethnographic research, are among the most commonly used to elicit user requirements as they offer the designer a rich nuanced understanding of user requirements [11]. These strategies depend highly on the ability of the designer to process the in-depth information gathered and translate it into technical requirements. These strategies are limited in their scalability due to the high cost in time and resources for each additional participant. Thus, these techniques work best when the selected participants are representative of the target user group. Lead-user theory and participatory co-design build on this work by integrating identified users into later stages of the design process [14, 15].

Matrix representations provide a framework for the designer to synthesize technical and qualitative information to provide a link between technical attributes and user requirements. There is a long history of improved product performance due to explicit requirement-attribute linking in Quality Function Deployment [18] and House of Quality [19]. However, both qualitative and matrix methods rely on the design team to synthesize nuanced information into engineering specifications. They are therefore susceptible to biases and misunderstandings inherent to human decision-making. Additionally, they require an upfront investment of resources to gather the raw qualitative data. They are best in situations in which the target user population is not well understood and are often used to identify new product opportunities. In response to these limitations, a large body of work has been developed which uses explicit representations of user preferences.

4.3.2 Explicit User Preferences

In translating user requirements into engineering specifications, a number of methods require the elicitation of explicit user preferences. For example, in Taguchi Robust Design and Design for Six-sigma meeting customer satisfaction is the objective of design decisions [20, 21]. In Decision-Based Design (DBD) the value of the designed artifact is maximized considering both the user and producer preferences [34]. Across all of these strategies, preferences are defined as the rank ordering by desirability of alternative outcomes to a decision [13]. Some methods further extend the explicit representation of preferences into a utility function, defined as a continuous measure of satisfaction an individual gets from the consumption of a good [97, 98, 99, 100]. A higher utility for a given design would indicate that the user would prefer that design to alternative designs with lower utility.

An established body of research has sought to improve the use of explicit preference information in product design. Discrete Choice Analysis (DCA) and Conjoint Analysis (CA) are the major demand analysis techniques used to capture user choice behavior. DCA relies on existing data sets of historical user choices to generate the probability an option is chosen based on available alternatives [22, 23]. CA uses survey responses of randomly generated ranked-choice questionnaires to generate a demand function based on product attributes [24, 25, 26]. Researchers have developed a number of modifications to DCA and CA to improve design outcomes including: using machine learning to improve the prediction accuracy [27], incorporating shape and aesthetic information in addition to technical attributes [28, 29], and incorporating the impact of sustainability preferences [30]. Recent work into Design Analytics (DA) has used data-driven machine learning techniques to improve the predictive power of the preference models [31]. Through cyber-enabled products, sensory devices, and the collection of usage data, DA can automatically identify product attributes which impact user choice behavior. All of these strategies seek to create an single accurate functional mapping between potential combinations of product features and user preference or utility.

Emerging work has looked into situations in which this preference function can change. Predictive trend mining [101, 102, 103] uses historical time-series data on user choices to make predictions about future user preferences. Research into evolvability defines a metric based on excess and modularity to estimate a system's ability to adapt to changing using preferences through reconfiguration [104]. These data-driven methods rely on a statistical analysis of surveys or historical observations to estimate the link between technical and user requirements. Although these strategies account for specific cases of changes in user preferences, the resulting models are sensitive to the quality and amount of data collected. These methods are best when collecting a large amount of data on user choices is feasible.

Making design decision by optimizing these explicit preference modelsis the basis of research into design optimization. Formal methods, such as Analytical Target Cascading and Decision-based Design, are based on incorporating an explicit representation of user preferences into the objective function of an optimization of the technical attributes of a product [32, 33]. This is in contrast to the more qualitative work on user needs in which a designer synthesizes gathered information into an understanding of user preferences. In optimization methods such as Multi-attribute Utility Theory user preferences are typically represented with utility, a measure of user preference constructed from possible configurations of product attributes [105, 33].

A utility function can be estimated using the techniques cited above including historical market data about user choices or survey methods such as conjoint analysis [106]. Modeling utility in this manner is limited in several ways. Survey respondents are asked to make hypothetical choices which may not accurately represent real-world choices well. Historical market data represents real choices but may not accurately represent preferences for products which are radically different from existing choices. Finally, all of these methods assume that each user has a latent utility function which can be modeled. Recent work has identified a number of ways in which this assumption is flawed, including demonstrating that the preference function is constructed based on available choices [107].

The estimated utility function is used in an optimization analysis to select a combination of product attributes or technical requirements which maximize the estimated user utility. A mismatch between the user preferences and the produced design can occur due to error in the estimate of the utility function through the mechanisms cited above. This mismatch can also occur if the user's utility function has changed between when the designer elicited the information and the product was produced. This study adapts work in intelligent control to propose an iterative design optimization process where the utility function and corresponding optimal technical requirements are updated based on user feedback throughout the life cycle of the product. The baseline comparison for this study is therefore a design optimization method in which the utility function is estimated at time zero and an associated optimal design is produced and used throughout the lifetime of the product.

4.3.3 Intelligent Control

Control methods are algorithms for minimizing a cost function by dynamically varying control parameters. Intelligent control is a category of control methodologies which is based on artificial intelligence. Intelligent control techniques have been developed to control highly-complex non-linear systems such as autonomous vehicles. Traditional methods, commonly based on differential equations, performed poorly on these types of tasks [108]. The intelligent control methods can be classified into machine learning, fuzzy logic, multi-agent system control, and metaheuristic algorithms such as bio-inspired and evolutionary optimization algorithms [108, 109]. This study is based on adapting a group of intelligent control methods based on Reinforcement learning (RL). Control schemes based on reinforcement learning, discussed in detail below, are currently based on objective measurements of the environment or users for feedback. This work extends this method to include subjective preferences as the feedback.

4.3.4 Machine Learning in Design

Many data-driven design methods use techniques from the three main machine learning paradigms (reinforcement learning, supervised learning, and unsupervised learning) to help guide design decisions. The majority of these strategies draw on either supervised learning to infer relationships between pre-identified sets of design variables and performance variables, or unsupervised learning to uncover patterns from unlabeled data without instructions [110]. For example, design Analytics tools and preference elicitation methods such as feature selection/engineering [27, 111], inferring utility (preference) functions [53, 33], and Conjoint based methods [53, 112] come from these two paradigms and require large existing data sets.

In contrast, reinforcement learning and active learning algorithms use information gathered through interactions with the user over time to make repeated decisions [113, 114]. Reinforcement learning techniques choose from a set of actions to maximize a reward given a certain state, while active learning techniques interactively query the user for data labels. These techniques can make repeated decisions based on past information and make different decisions as new information is collected. This is a potentially useful characteristic in situations where it is impossible to gather preference data for all the design points in advance as the product users and usage contexts are unknown. Furthermore, it may not be feasible to explore all the design possibilities to generate the labeled data point for the supervised learning methods, and the exploration should be led by the feedback from the user in a way that the trade off between exploration and exploitation is balanced. Adding to that complexity is the possibility of dynamic user preferences where the user preference evolves and changes over time [101, 103]. Under these conditions, a product which learns how to respond and react as it interacts with the user over time would be beneficial.

4.3.5 Research Gap

Static user preferences are a key underlying assumption of product design. However, previous work cited above and initial data from the illustrative case study indicate that user preferences can vary greatly on the time scale of the lifetime of the product. A product may still be functional but not generate utility for the user. Existing design methods are not capable of producing systems which change their functionality in response to changing user preferences. This study seeks to extend current design methods to include dynamic preferences by adapting reinforcement learning algorithms for use in designing physical systems. If successful, the design for dynamic preferences framework will enable the creation of a new class of "smart" physical devices which adapt to changing user preferences. Given this context, this study seeks to answer the following research questions:

1. What is the impact on user preferences over the lifetime of a product of a static design strategy versus a dynamic strategy based on reinforcement learning?

2. Under what conditions do different reinforcement learning algorithms impact user preferences over the lifetime of a product?

In answering these questions, the illustrative case study presents simulated user preference data. This was done to compare static and dynamic design strategies under a wide variety of conditions and assumptions. Further work is needed to identify which conditions and assumptions are valid for real use cases.

4.4 Design for Dynamic Preferences Framework

The following design framework is developed under the assumption that the user's preference function will change over time. To illustrate the performance of the new framework the following illustrative case is proposed.

4.4.1 Illustrative Case: Design of a Variable Stiffness Prosthetic Ankle (VSPA)

This study considers the design of a variable stiffness prosthetic ankle (VSPA) [115]. Figure 4.1 shows an isometric view of the VSPA. The VSPA has two advantages over traditional ankle prostheses: 1) it has a nonlinear, custom torque-angle curve that can more accurately mimic biological ankle mechanics, and 2) it is able to perform online modulation of the overall stiffness for different mobility tasks. Figure 4.2 shows the different torque-angle curves obtained experimentally by varying a stiffness control parameter, x. The stiffness control parameter represents the position of a mechanical slider on the VSPA and varies between zero and one hundred percent of the allowable range. The design intent of this prosthesis is to enable variation in the torque-angle curve and ankle response during different mobility tasks, such as moving from walking on flat ground to walking up stairs.

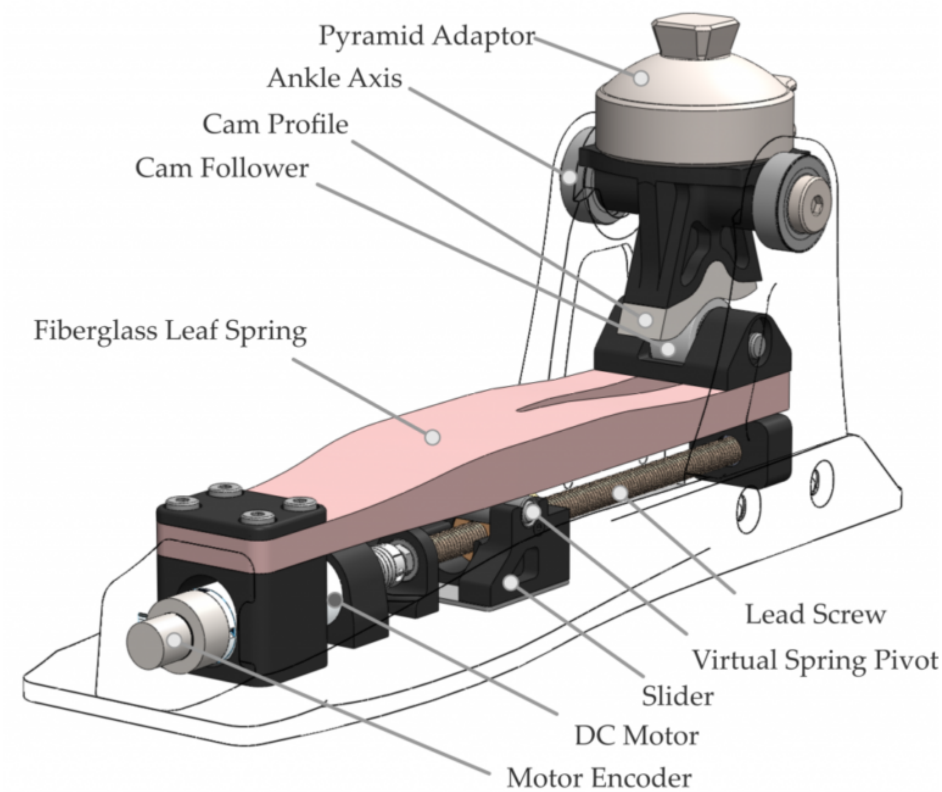

Figure 4.1: Isometric view of variable stiffness prosthetic ankle (VSPA).

In this situation, there is a single discrete design variable, x, which has possible values: $x \in \Omega = [0, 100]$, where Ω is defined as the design space or all feasible values of x. The stiffness control parameter can be adjusted to suit the preference of the user. Currently, users visit a technician to be fitted with the prosthesis. During this visit, the user walks on a treadmill for 5 minutes and adjusts x until settling on an optimal stiffness parameter x_* with respect to their preference [115].

This is a repeatable process with users arriving at the same optimal stiffness during laboratory testing [115]. However, data collected from the same individuals tested a year apart shows a variety of behaviors with regards to optimal stiffness parameter. Figure 4.3 shows the preferred x_* for three individuals measured a year apart for the same mobility task. The preferred x_* remained the same for one patient and was drastically changed in another. This variance could be due to a number of factors, including changes in body weight, soft tissue mass, and bone structure over the course

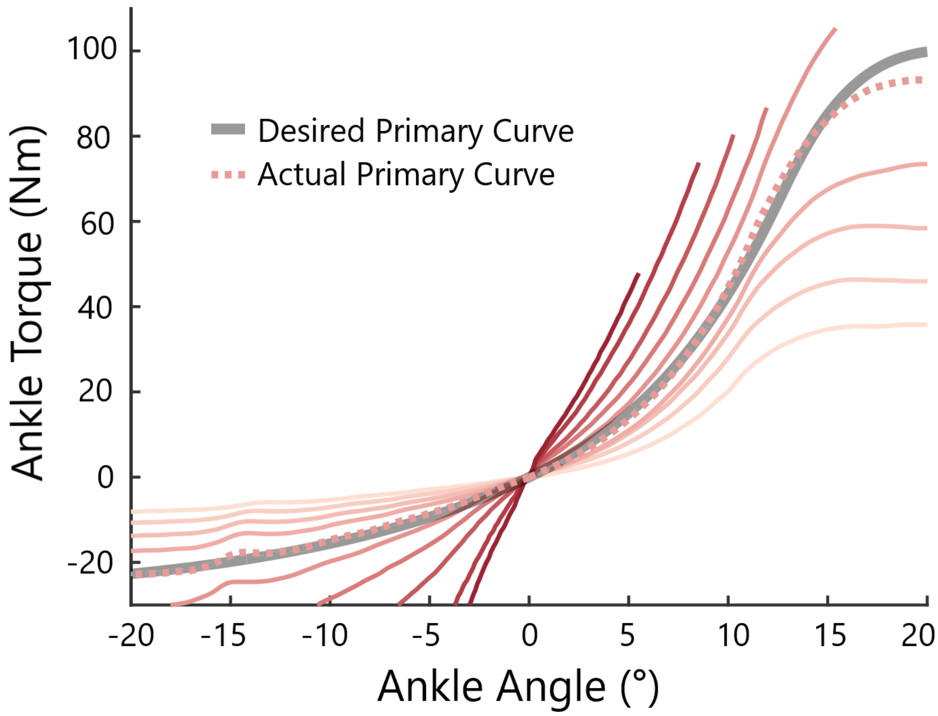

Figure 4.2: Torque-angle curves for different stiffness parameters x. For lightest curve $x = 0$, for darkest curve, $x = 100$.

of the year. The proposed design problem is then how to choose x throughout the lifetime of the prosthesis.

4.4.2 Model Assumptions

In order to illustrate the utility of the proposed framework a number of simplifying assumptions have been made. In this case, user preference is modeled as a function of the stiffness parameter, $P = f(x)$. The preference function is assumed to be bell shaped and scaled to have a maximum of 1. Time is limited to a year period with an evaluation of the user's preference on a daily basis, $t \in [0, 365]$. At each time point, it is assumed that the preference value of 1 is achievable for all individuals. A result of this is that there exists an optimal point $x_* \in \Omega$ such that $P(x_*) = 1$. Previous work on the physical design of the prosthesis suggest that users can accurately distinguish between settings of x that are greater than eight percent apart [116]. Thus the design problem is formulated as a discrete selection of $x \in [0, 10, 20, 30, 40, 50, 60, 70, 80, 90, 100]$. This reduces the RL formulation to a k-armed bandit problem with eleven choices.

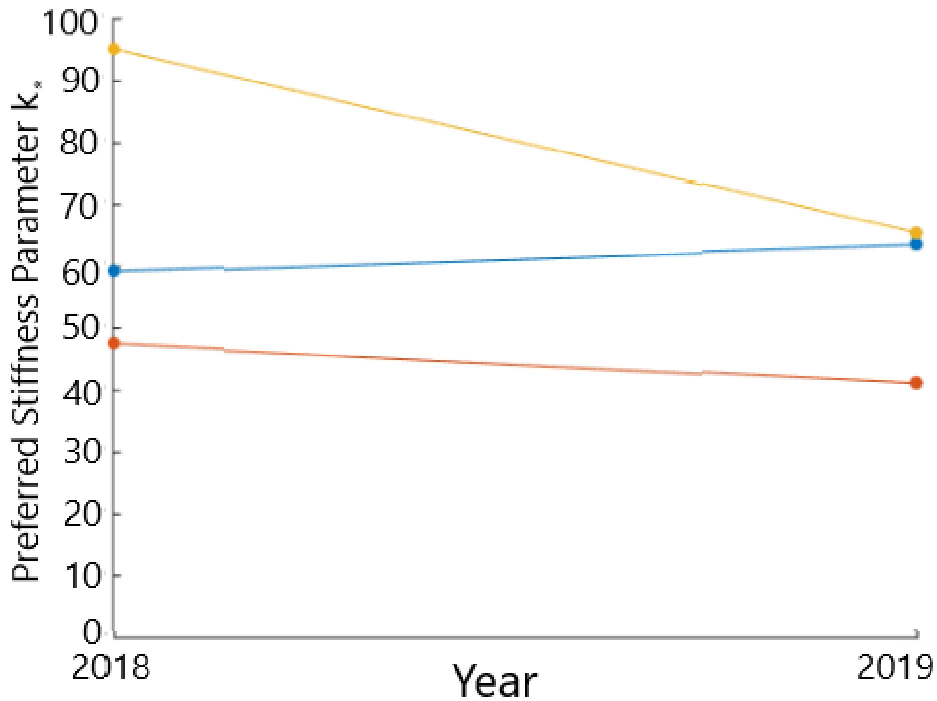

Figure 4.3: Preferred stiffness parameter x_* from 2018-2019.

4.4.3 Metrics

Given these assumptions, it is possible to calculate several performance metrics for any given design strategy. User preference is evaluated at each time point, t, for $t \in [0, 365]$. The first metric is the cumulative sum of preference values over the time period.

$$\begin{aligned} C &= \sum_{t=0}^{365} P(t) \\ C &\in [0, 365] \end{aligned} \tag{4.1}$$

The second metric is the average of preferences up until any time t (cumulative mean).

$$\begin{aligned} \bar{P} &= \frac{\sum_{t=0}^{j} P(t)}{j} \\ \bar{P} &\in [0, 1] \end{aligned} \tag{4.2}$$

Under both of these metrics, a higher score represents higher performance. However, C can only be measured across the whole time period, while $\bar{P}$ can be evaluated at any time t.

The final metric compares a design strategy to the baseline static strategy. The crossover point, t_c, is the minimum time t at which the average preference, $\bar{P}(t)$, for the selected RL strategy is higher than the average preference for the baseline static strategy.

$$
\begin{aligned}
& t_c = \min_t \bar{P}_{static}(t) < \bar{P}_{RL}(t) \\
& t_c \in [0, 365]
\end{aligned}
\tag{4.3}
$$

If the average preference of the RL strategy never exceeds the static strategy, t_c is set at 365.

4.4.4 Design Strategies

Starting from the simplest case, we try to control the VSPA foot using reinforcement learning algorithm for the special case where only one state is involved. Under this assumption, the decisions made by the system do not change the state of the environment. This type of problems can be handled using a set of methods called k-armed bandit problems.

4.4.4.1 K-Armed Bandit Problem

K-armed bandit problems are a group of problems dealing with the question of optimally and repeatedly choosing from k alternatives or actions each with a hidden reward distribution. The goal is maximizing the expected total reward over some time period. There are several challenges that are faced by algorithms designed for this type of problems. There are no prior knowledge on what choice yields the highest reward and the algorithms should figure it out over time. Moreover, there is always a trade off between exploration and exploitation. In other words, the algorithm can settle in early on with the best choice based on some early estimation and risk losing on some unexplored choices with higher rewards. On the other hand, exploring may result in hitting suboptimal choices more often which in turn results in suboptimal total reward [113].

Translating the case study to a k-armed bandit problem, the rewards are the preference ratings fed to the system by the user on each day. The k arms or actions are the set of 11 values for x. The objective is to maximize the total (cumulative) preference over a year of usage.

Four algorithms for deciding x are tested in this illustrative case. The baseline is a static strategy under which x_* is determined at $t = 0$ and remains at that setting for the entire time period. This models an initial visit to the technician and calibration of the prosthesis to maximize the user's preference at that time. Three reinforcement learning algorithms are also considered for deciding x at any time t: 1) an ϵ-greedy algorithm, 2) an Upper-Confidence Bound (UCB) algorithm, and 3) a dynamic reinforcement learning algorithm which is an ϵ-greedy algorithm that weights recent evaluations more highly than older ones. All of the RL strategies have no information about x_* at time $t = 0$, but use information from previous evaluations to decide x. The initial x for these strategies is chosen randomly from a uniform distribution.

4.4.4.2 ϵ-Greedy

In all of the RL strategies, the actions are ranked based on some measures updated at time t, and the action ranked top is selected. If the measure is defined as the average reward of action a up until time t, the algorithm is called greedy. $Q_t(a)$ denotes the ranking measure for greedy algorithm.

$$Q_t(a) = \frac{\text{sum of rewards when } a \text{ taken prior to t}}{\text{number of times } a \text{ taken prior to t}} \tag{4.4}$$

The action A_t is then chosen as

$$A_t = argmax_a Q_t(a) \tag{4.5}$$

As the name suggests, greedy algorithm focuses more on exploitation and suffers from exploration issues. Due to the fact that always the action with highest Q is picked, the chances for the possibly better actions that performed poorly at the beginning would be zero as the algorithm will only pick the same action that performed better in the early stages. To handle this issue a new parameter ϵ is introduced. After each ranking of the actions, with a small probability ϵ, the algorithm dismisses the ranking and picks the action A_t randomly from all the actions with equal probability. Although the ϵ-greedy method behaves same as greedy most of the time, it ensure exploration of all the actions given a sufficiently long time span.

4.4.4.3 Upper-Confidence Bound (UCB)

Upper-Confidence Bound (UCB) algorithm accounts for uncertainty around an action as an advantage in the rankings by $Q_t(a)$. The rationale behind this approach is that the more uncertain the action, the higher chance of it being a better one (or a worse one), hence resolving the greedy algorithm exploration issues by exploring the actions with higher information content. The ranking measure for UCB is defined as

$$Q_t(a) + c\sqrt{\frac{\ln t}{N_t(a)}} \tag{4.6}$$

Where $N_t(a)$ denotes the number of times that action a has been chosen prior to time t. $c > 0$ sets the degree of exploration.

4.4.4.4 Dynamic ϵ-greedy

$Q_t(a)$ measure introduced before is a simple average over all the previous rewards of a, $R_t(a)$, weighting them equally. Under the circumstances of evolving and changing reward distributions, the simple average seems inadequate as more recent rewards give more information as to which action is the optimal one. One way of doing the weighted averaging is through the step-size parameter $\alpha \in (0, 1]$ with values closer to 1 giving more weight to recent rewards. For a specific action, a, with initial $Q_t(a)$, Q_1, the weighted reward average can be calculated recursively as

$$Q_{n+1} = Q_n + \alpha(R_n - Q_n) \tag{4.7}$$

Note that the index n here defines the place of the variables among the set of previous Q and R for a specific action and it is different from t which is the time index.

The action A_t is selected as

$$A_t = argmax_a Q_t(a) \tag{4.8}$$

4.4.5 Preference Function Construction

Performance of a design strategy is considered across a range of constructed preference functions. Preference functions are defined by three characteristics. First, the initial optimal stiffness parameter $x_{*,i}$, x_* at $t = 0$ and final optimal stiffness parameter $x_{*,f}$, x_* at $t = 365$ are chosen to be in $\Omega = [0, 100]$. Second, the preference function is defined as either deterministic or stochastic. A deterministic preference function, $P(x, t)$, is defined directly from x and t. A stochastic preference function is defined as a normally distributed random variable whose mean is the deterministic preference function. Finally, the extrapolation of the constructed function is either linear, exponential, or periodic. In each case x_* at t is a linear, exponential, or periodic extrapolation between the initial optimum $x_{*,i}$, and the final optimum $x_{*,f}$.

In general, at time t the deterministic preference function is

$$P(x, t) = e^{-\frac{(x - x_*(t))^2}{2\sigma^2}} \tag{4.9}$$

Where $x_*(t)$ is the optimal stiffness parameter at t and σ sets the width of the preference function. Any function with the behavior of peaking at 1 at some single input and approaching zero as the input gets further from the optimal point x_* would yield similar final results. However, Gaussian was chosen so that 1) the preference value is never negative no matter how far from x_*, 2) the behavior of the function can be simply managed by two values $\mu = x_*$ and σ (preference sensitivity), 3) preference value decreases monotonically as distance between x and optimal stiffness

Table 4.1: Extrapolation functions, $x_*(t)$.

Extrapolation	$x_*(t)$
Linear	$x_{*,i} + \frac{(t-t^i)(x_{*,f}-x_{*,i})}{t_f-t_i}$
Exponential	$x_{*,f} + (x_{*,i} - x_{*,f})e^{-\lambda(t-t_i)}$
Periodic	$x_{*,i} + \frac{(t-t_i)(x_{*,f}-x_{*,i})}{t_f-t_i} + A\sin(2\pi f \frac{t-t_i}{t_f-t_i})$

increases. Table 4.1 shows the three extrapolations of x_* as a function of t for linear, exponential, and periodic optimal preference changes.

In this table, t_f is the final time, t_i is the initial time, A is the amplitude, and f is the frequency of the periodic change (cycle per time span, in this case cycle/year). x_* out of interval $\Omega = [0, 100]$ would be set to the value of nearest bound.

4.4.6 Simulation Characteristics

To characterize the performance of the design strategies under different conditions, 10000 preference functions are constructed for each function type : 1) deterministic linear, 2) deterministic exponential, 3)deterministic periodic, 4) stochastic linear, 5) stochastic exponential, and 6) stochastic periodic. For each constructed preference function, $x_{*,f}$ and $x_{*,i}$ are chosen randomly from a uniform distribution of $\Omega = [0, 100]$.

At time $t = 0$, the static strategy chooses an initial x from $x = [0, 10, 20, 30, 40, 50, 60, 70, 80, 90, 100]$ which is closest to $x_{*,i}$ for the constructed preference function. This x will remain the selection for the entire time period. The RL strategies randomly choose an initial x from the same set with no knowledge of $x_{*,i}$. The preference function, $P(x, t)$ is evaluated at $t = 0$. The RL strategies then choose a x for time $t = t + 1$ based on previous selections of x, evaluations of $P(x, t)$, and their decision policies. The three performance metrics, C,$\bar{P}$, and t_c are then calculated for each design strategy.

4.5 Results

As an example, consider a deterministic linear preference function constructed with $x_{*,f} = 70$ and $x_{*,i} = 40$. Figure 4.4 shows the preference with respect to time and stiffness parameter x. Figure 4.5 shows average preference $\bar{P}$ over $t = [0, 365]$ for this preference function for all four strategies.

Table **??** shows the mean of the cumulative sum C for each design strategy over the 10000 preference functions. The metric $C \in [0, 365]$ with $C = 365$ representing the maximization of user preference on every day in the year period. In a two-sample t-test with the mean of the static

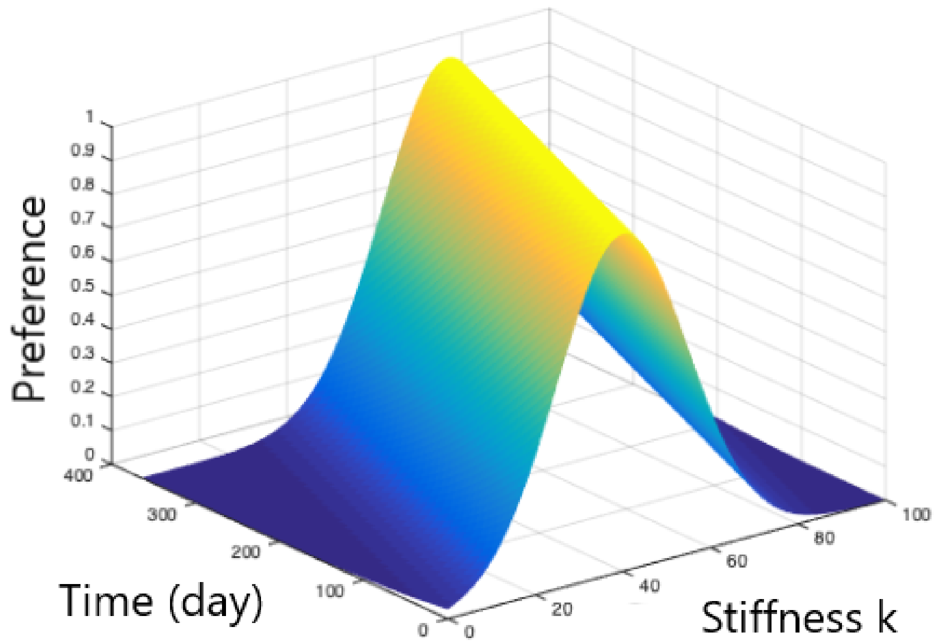

Figure 4.4: Linear preference function with $x_{*,I} = 40$ and $x_{*,f} = 70$.

Table 4.2: Average cumulative sum of preference c.

	Deterministic			Stochastic		
Strategy	Lin.	Exp.	Per.	Lin.	Exp.	Per.
Static	222.54	167.46	125.58	215.47	168.97	133.14
ϵ - greedy	262.04	277.87	136.67	253.19	266.35	144.37
UCB	315.19	316.97	172.54	300.13	301.86	188.82
Dyanmic RL	279.63	282.73	145.94	273.59	275.84	152.62

design strategy, the mean of each RL strategy was statistically significantly different with a p-value indistinguishable from 0.

Figure 4.6 shows box plots for the cumulative preference C for each design strategy for the deterministic preference functions. As seen in the above tables, relative results from the deterministic and stochastic preference functions were similar. Thus, for space considerations only the deterministic results are shown here.

4.6 Discussion

One notable result from this study is that reinforcement learning strategies can outperform the static design strategy in cases where preferences change even with no prior information. Although expected in the construction of the case, as seen in Table 4.2 the size of the effect is large over

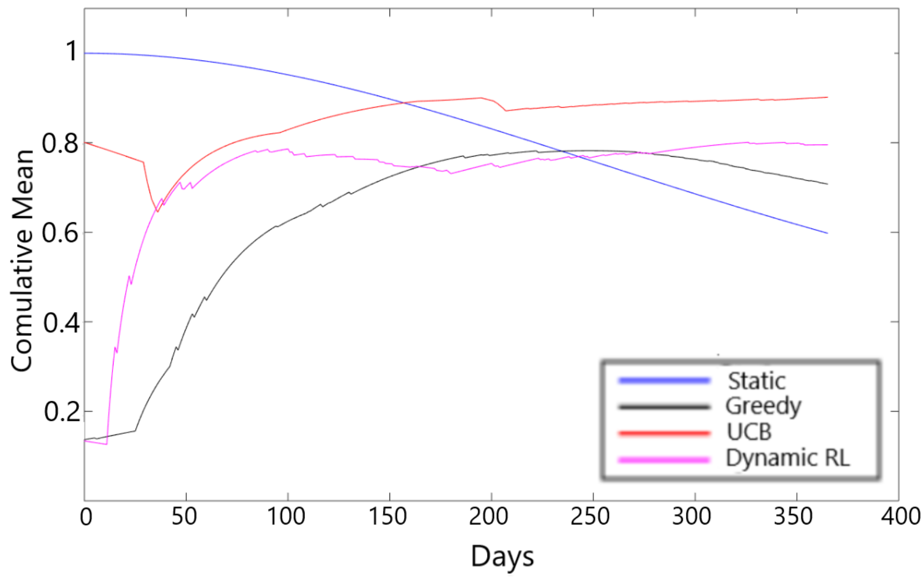

Figure 4.5: Average preference up until t_i, $\bar{p}$ for four design strategies (linear, $x_{*,I} = 40$, $x_{*,f} = 70$.)

Table 4.3: Average crossover time t_c.

	Deterministic			Stochastic		
Strategy	Lin.	Exp.	Per.	Lin.	Exp.	Per.
ϵ - greedy	235	217	162	151	88	80
UCB	190	180	122	122	55	47
Dynamic	236	213	164	152	73	67

the range of possible preference functions. Reinforcement learning strategies perform even better when the preference functions change quickly as in the exponential case. Given the large cost associated with current preference modeling tasks which lead to a static outcome and evidence suggesting preferences are dynamic, this result suggests that further work into understanding how to implement reinforcement learning on physical products may produce significant cost savings and performance gains in comparison to existing techniques.

Another notable result is that the UCB strategy appears to perform best across a wide range of conditions. This is in agreement with some of the work in machine learning [113]. On average across all sampled preference functions, the cross over date for UCB is lower and the cumulative sum metric is higher. This result suggests that UCB should be used in cases where nothing is known about the dynamic nature of the preference function. The other reinforcement learning may be used depending on the structure of the preference function. For example a dynamic reinforcement learning strategy should be used for certain periodic preference functions.

The results presented in Table 4.3 demonstrate that because of the difference in knowledge

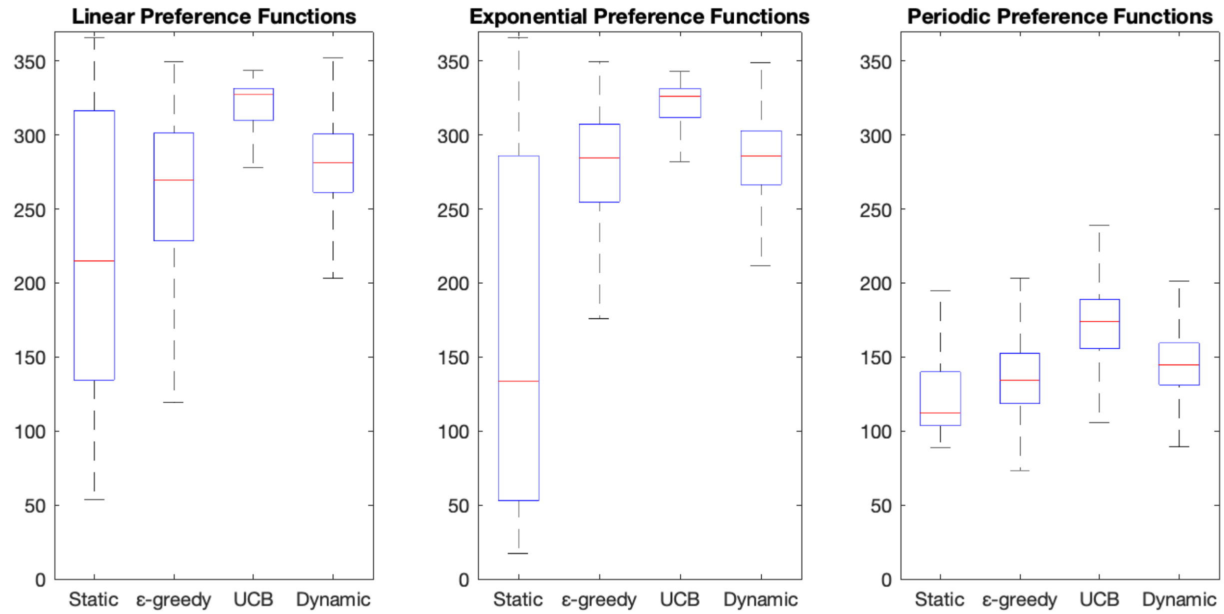

Figure 4.6: Cumulative preference for deterministic preference functions.

at $t = 0$, the static design is preferred initially and there is a significant time period before the reinforcement learning techniques surpass the static design. This is due to the assumption that the user preferences will not change quickly. This result suggests that a hybrid approach where the reinforcement learning technique uses some knowledge of the initial preferences would improve user preferences over the lifetime of the product. Further work is need to establish the trade-off between cost of eliciting user preferences and improvement in the RL strategy performance.

Finally, the presented results represent sampling from a population with extreme heterogeneity in their preference functions. The ten thousand samples are drawn randomly from a possible 10201 preference functions for each shape. Thus, each sample likely has a different $(x_{*,i}, x_{*,f})$. The reinforcement learning techniques outperformed the static strategy because the majority of these preference functions were dynamic, $x_{*,i} \neq x_{*,f}$. However, the assumption that the static strategy chose $x_{*,i}$ such that $P(x_{*,i}) = 1$ is a generous assumption for the static case. In the illustrative case, this assumption represents the initial visit to the technician and a perfect calibration to the individual user. For most products, market segmentation is used to try and identify a user population with similar preferences. The preference elicitation techniques are used to estimate a single preference function for the entire population and a single mass-produced product is designed for that population. Thus, a static strategy will likely not maximize user preference even at the initial time point when the most information is known. The reinforcement learning strategies have the ability to "personalize" the functionality over time and may outperform the static strategy even in cases where the preference function is not dynamic.

This study is limited by several factors. First, although grounded in the observed empirical findings from previous work, the data is simulated and the simplifying assumptions may not be true for actual user populations. The preference function may not reach the maximum utility at every time point and may not be Gaussian. Changes to these assumptions may affect the size of the impact of the reinforcement learning techniques. This study is also based on the assumption that the user is able to evaluate their preference with a ranking between 0 and 1. Previous work has demonstrated that other feedback mechanisms, such as choice-based methods, may be more accurate in eliciting user preference evaluations [53]. Using a different feedback mechanism could influence the performance and selection of RL strategies. Finally, although the preference function was continuous, a discrete version of the k-bandit problem was used to model the physical system of the illustrative case. Other design tasks may involve continuous design variables. Continuous variable formulations of reinforcement learning algorithms used here exist but their performance may not be accurately represented by the results of this study.

4.7 Conclusion & Future Work

This study presents results from an illustrative case examining the use of a Design for Dynamic Preferences framework on the design of a Variable Stiffness Prosthetic Ankle. Simulated data across a wide range of preference function conditions suggests that reinforcement learning strategies could provide advantages over existing design methods. The three reinforcement learning strategies tested could improve user preferences over the the lifetime of the product when compared to static strategies without the upfront cost of user preference elicitation. In particular, this study suggests that a UCB strategy may be the best if there is no knowledge about the preference function shape ahead of time. Performance of dynamic methods could improve if initial information is incorporated. In conclusion, this study lays out a road map for enabling the design of a new class of "smart" products which use feedback from the user to optimize product features for changing user preferences.

This study suggests many additional avenues for research into Design for Dynamic Preferences. Future work will include testing different feedback mechanisms such as rating/ranking systems versus choice-based systems. This will improve the accuracy of preference evaluations in real systems. Future work will also include testing of reinforcement learning strategies under different conditions which relax the simplifying assumptions made in this study. For example, design problems with continuous design variables will be investigated. The utility of this framework on the personalization of products for a population will also be studied.

CHAPTER 5

Understanding User Willingness to Interact with Adaptive Engineered Systems

This chapter was coauthored with Jesse Austin-Breneman.

5.1 Abstract

Adaptive engineered systems which are responsive to dynamic user preferences require a constant flow of preference feedback from the user to operate effectively. However, prior work in human factors and survey fatigue suggests a number of factors can lead to lower quality data or non-response. An empirical pilot study of adaptive office chair was conducted to investigate factors affecting user willingness to repeatedly interact with the adaptive system. A statistical analysis of the results shows that desirability of the system state impacts the reported user willingness to interact over long periods of time. Results also found no significant difference between responsive system state and an unresponsive desirable state. These findings lay out a road map for future designers of these systems to tailor the characteristics of the reinforcement learning algorithm to maximize user willingness.

5.2 Introduction

The rise of inexpensive computation and communication technologies has led to an explosion in smart devices [1]. From smart refrigerators to intelligent medical devices, an increasing number of engineered systems are able to make data-driven decisions automatically by leveraging large collected data sets [117]. On-board computation, sensors, and access to shared data sets have enabled systems to tailor their functionality to individual user's needs. For example, smart refrigerators are able to detect missing ingredients, construct personalized shopping lists, and share it with your phone [118]. These adaptive hardware systems are defined as systems which modify their functionality in response to environmental changes [119]. A major challenge for designers of adaptive systems is to predict and design for future systems states in an exhaustive manner. The number of potential alternatives are numerous by construction and non-linear emergent behavior arising from interactions between a dynamic environment and the system are difficult to model. This is especially true for large adaptive systems, such as fleets of autonomous vehicles where the number of interactions can scale exponentially. Recent work in computing and information technology has demonstrated the potential of these types of reconfigurable technologies to become self-integrating and self-adaptive in order to mitigate these uncertainties [120]. A self-adaptive embedded system may be more fault-tolerant and be able to react more quickly to collected information across a wider variety of operating conditions [121]. Although there is a rich body of literature exploring these issues from a computational perspective, these systems are currently limited to responding to environmental changes. Prior work by the authors has identified potential benefits to creating adaptive engineered systems which respond to subjective user preferences in addition to environmental changes [54]. A smarter refrigerator would be able to not only detect missing ingredients but construct shopping lists based on an understanding of what types of food you like.

In working towards the design of adaptive engineered systems which respond to dynamic user preferences and dynamic environments, Arezoomend, et al. [54] proposed incorporating reinforcement learning (RL) into a design optimization framework. In this new method, RL algorithms are used to constantly update product attributes as the user interacts with the product. Similar to dynamic software, this type of adaptive engineered system would use user-provided feedback on their subjective preferences to update an objective function for setting product attributes. For example, a smart knee brace could collect both objective bio-mechanical data and subjective comfort information from the user in order to set its stiffness parameters. Simulations of a variable stiffness prosthetic ankle showed that the use of RL algorithms to control the stiffness parameter in this fashion could improve the user satisfaction over the lifetime of the prosthesis under certain conditions.

Because RL algorithms perform better with more observations, a constant flow of information is critical to the functionality of an adaptive engineered system of this type. Thus, designers of adaptive engineered systems need to ensure that users will continue to provide this data to the system over its entire lifetime in order to have successful outcomes. In similar situations, such as in-app surveys, prior work has reported significant levels of non-response [122]. This work therefore focuses on understanding the factors which impact a user's choice to continue interacting with an adaptive engineered system to provide data. Specifically, this study seeks to determine if there are characteristics of system itself which could be manipulated by designers to maximize user willingness to interact. Results could then be used to adjust RL algorithm behavior to be more suitable for this application. An empirical pilot study examining an adaptive office chair is presented to test various factors.

5.3 Related Work

This paper works builds upon work in a number of fields including design research, human factors, and machine learning.

5.3.1 Preference Elicitation Method Evaluation

The phenomena at the center of this study is a preference elicitation task. The adaptive engineered system is asking the user to evaluate the desirability of the current system state. Preference elicitation methods have been well investigated by different disciplines from computer science [123] and human computer interaction [124] to marketing [125], healthcare policy [126], decision science [127] and economics [128]. In their systematic review of preference elicitation methods, Ryan, et al. list five criteria for evaluating quantitative elicitation methods, such as the type required for RL algorithms [126]. The five criteria are validity, reproducibility, internal consistency, acceptability to respondents, and cost. Because this study is focused on user willingness to interact, acceptability is the most relevant performance metric. Acceptability has been previously operationalized using time to complete, response rates, and completion rates[126]. This work draws upon this literature to help define design objectives for the RL algorithm.

5.3.2 Survey Fatigue

In this work, the impact of aspects of RL algorithm behavior on user willingness for providing ongoing feedback will be studied. One major contributor to nonresponse rates found in disciplines which use survey-based techniques is survey fatigue [129]. Survey fatigue is defined as the situation in which respondents become bored or tired of answering questions during the administration of the survey [130]. RL data collection is most similar to survey fatigue for panel surveys as they involve repeated survey iterations over a time period [131]. Survey fatigue can lead to poor quality data in a number of ways. Respondents experiencing survey fatigue have been shown to skip questions, quickly try to answer questions without really considering their answer, and provide random responses [132].

In seeking to mitigate survey fatigue through careful questionnaire design, researchers have identified factors that are known to influence survey fatigue. In general, survey fatigue is minimized when questionnaires are designed to balance the respondent's motivation to take the survey with the burden imposed by taking the survey [133]. Factors which affect motivation can include compensation, intrinsic interest, and potential impact. Respondents who believe their voice will be heard and have an effect on the subject of the survey are more likely to fully complete the survey [134]. Factors which affect the burden imposed by the survey can include survey length, complexity, and question type [132]. Survey length contributes to a time burden, cited by many nonrespondents as the determining factor for nonresponse [135]. More complex surveys and certain question types may impose a greater cognitive burden on the respondent and lead to greater fatigue. Of most relevance to this work, repeated short surveys have also been found to induce survey fatigue [135]. Although these results are specific to survey questionnaires, this study uses these concepts as potential factors impacting user willingness to interact with adaptive engineered systems.

5.3.3 Human Factors of Adaptive Systems

Research in human factors has long considered factors affecting user interactions with adaptive systems. One relevant area of research is trust in an adaptive systems [136]. Prior work has found that users generally have lower levels of trust for more adaptive systems [137]. System transparency, user control, perceived responsiveness, and perceived competence are all factors which can strengthen user trust [138]. Trust as studied in this literature relates to the users confidence in the performance and credibility of the adaptive system. This may translate to increased willingness to continuously interact and provide feedback. This paper draws upon this work to identify potential factors for study.

Because this study focuses on collecting subjective user preferences, the exemplar adaptive system should have clear subjective performance measures. Additionally, optimal configurations should vary across individuals with no easily correlated objective measure. Designing a comfortable office chair was chosen for the case study as it meets all of these conditions and is a well-studied problem in human factors research. There are a number of design heuristics for creating good office chairs, but predicting the optimal configuration with respect to comfort for a given individual remains extremely difficult [139]. Based on prior work testing office chairs, critical dimensions for improving comfort include seat thickness, backrest angle, and location of lumber support [140]. This paper uses experiments testing office chair comfort as the basis for testing user willingness to interact.

5.3.4 Research Gap

Prior work by the authors proposed using reinforcement learning algorithms to automatically control adaptive engineered systems. These algorithms require a constant flow of preference information from the user to function correctly. This study seeks to fill a gap in the scientific understanding of factors affecting user willingness to interact with adaptive systems. Because factors affecting the burden posed by the interaction are highly context-dependent and not necessarily within the control of the designer, this study focuses on factors affecting motivation. Prior work in survey fatigue and human factors have identified several promising factors, including perceived competence and system responsiveness. Given this background, the researchers seek to answer the following research questions:

1. Does user preference for the current system state impact user willingness to continue interacting with an adaptive system?

2. Does user perception of system responsiveness impact user willingness to continue interacting with an adaptive system?

To answer these questions, an empirical study of user willingness to interact with an adaptive office chair was conducted. Statistical analysis of results are presented to explore the impact of system state, system responsiveness, and interactions between these factors.

5.4 Methodology

To test the hypotheses, a set of human experiments was conducted using an adjustable office chair as adaptive system. The experimental design described below closely follows the setup found in Groenesteijn, et al. [141]. Three major modes of design space exploration/exploitation behavior can be imagined for the reinforcement algorithms based on the work done in chapter 4, and the literature on human factors and survey fatigue. Usually the RL algorithms start with a rapid design space exploration followed by an exploitation phase of the desired state. If no desirable state exists in the design space, we can imagine a constant undesirable state for the system. Therefore, the three treatment scenarios explored in this study are 1) comfortable configuration, 2) uncomfortable configuration, and 3) changing configuration where the first two are close to the concept of desirability and the last is similar to responsiveness discussed in human factors and survey fatigue literature. Participants explored all three scenarios in a random order and their willingness to provide feedback was measured through a survey mechanism. User willingness was measured on a Likert scale using a survey after all the scenarios were explored by the subject.

5.4.1 Hardware

Three adjustments are available on the office chair, namely backrest lock angle, lumbar support pillow height, and seat cushioning. Figure 5.1 shows the backrest angle. A digital angle finder was attached to the backrest to accurately measure the backrest angle throughout the experiment.

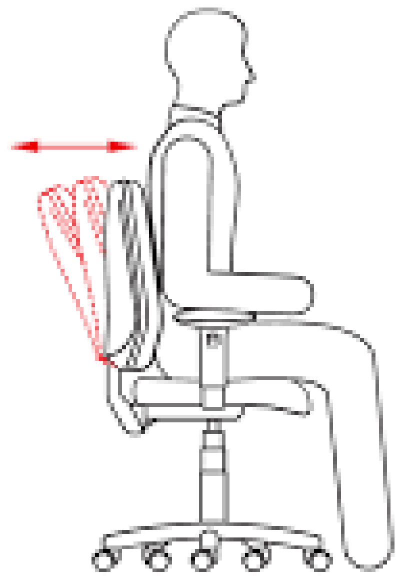

Figure 5.1: Backrest lock angle.

A cylindrical pillow was used as the lumbar support, fig. 5.2. The pillow can be placed at different heights on the backrest using VELCRO straps.

A cushion is used as the final adjustment, fig. 5.3. Unlike the other two adjustments that are continuous, the cushion is either used or not used. Figure 5.4 shows the setup.

5.4.2 Procedure

The subject population was recruited through email consisting of 29 students from the University of Michigan department of Mechanical Engineering. The experiment started with a warm-up session. In the warm-up session, the subjects were asked to explore different adjustments to get the feel of the configuration space. They were then asked to provide two configurations, one that they felt very comfortable in and one that they felt extreme discomfort in. The two configurations along with their rating were recorded for each participant. The ratings ranged from -5 to 5 with -5 being extreme discomfort, and 0 and 5 being neutral and very comfortable, respectively.

The three scenarios for each participant were ordered randomly each taking approximately 20 minutes. During the experiment sessions the subjects were asked to do simple tasks on a desktop PC while sitting on the chair. The tasks included typing a set of paragraphs and watching videos.

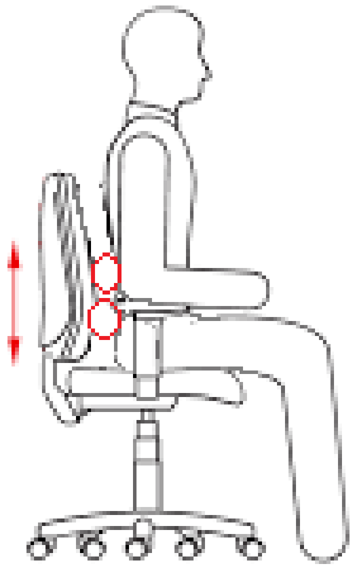

Figure 5.2: Lumbar support height.

The subjects provided feedback on the chair comfort every minute using the same -5 to 5 rating system. After each rating depending on the session scenario a new chair configuration was set and the process was repeated 15 times for each scenarios with the comfortable and uncomfortable scenarios using the configurations recorded in the warm-up session. The subjects could skip each session at anytime during the experiment. Participants were compensated with a $25 gift card regardless of how many sessions they skipped and how many ratings they provided.

The experiments were concluded with a final survey to rate their willingness to provide feedback for each scenario on a five-point Likert scale. Their willingness ratings were used to test the hypotheses.

5.5 Results

Upon finishing the three experiment sessions, the subjects were given a survey asking their willingness for providing prolonged feedback for each scenario. The willingness measures are converted to numbers with *Unwilling*, *Somewhat Unwilling*, *Neutral*, *Somewhat Willing*, and *Willing* converted to 1, 2, 3, 4, and 5, respectively. There has been a longstanding dispute about the most valid way to handle Likert scale data [142]. Parametric methods such as T-tests and Analysis of Variance (ANOVA) assume a continuous and normal distribution for the data, while the Likert data is discrete, ordinal, and with a limited range. Therefore, non-parametric methods such as Mann-

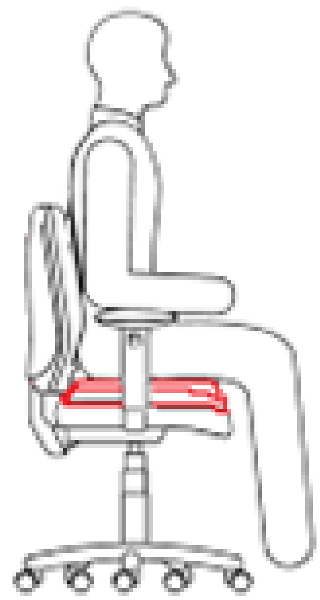

Figure 5.3: Seat cushioning.

Whitney test [143] with no assumptions on the underlying distributions are advised to be used. However, many studies have shown that generally the two sets of methods give in almost the same significant and insignificant results [144, 145, 146]. Throughout this study, the Likert data will be analyzed using parametric methods such as ANOVA.

A total of 29 subjects participated in the experiment, producing 29 ratings on a 1 to 5 scale for each scenario. Table 5.1 shows the summary of the rating data. The comfortable and changing scenarios have very close mean and variances, while the uncomfortable scenario has a smaller mean and a larger variance. Throughout this section the differences between the treatment groups will be investigated rigorously using statistical tools.

Treatment	*Count*	*Mean*	*Variance*
Comfortable	29	4.37	0.88
Uncomfortable	29	3.62	1.60
Changing	29	4.37	0.81

Table 5.1: Experiment data summary.

5.5.1 One-Way ANOVA

Before getting into the differences among each pair of groups and quantifying the impact of each treatment, a one-way ANAOVA is preformed on the data. One-way or single factor ANOVA is

Figure 5.4: Hardware setup.

used to analyze the differences among means of populations that are different only in one factor [147]. if ANOVA fails to reject the null hypothesis, usually no further comparison of population means is necessary. In this case, the only factor is the treatment, which can be either comfortable, uncomfortable, or changing. ANOVA shows whether at least the mean of one treatment group is different from the others'. More specifically, the one-way ANOVA tests the following hypotheses:

- H_0: All treatment group means are equal.
- H_a: At least one group mean is different from the rest.

Results of one-way ANOVA for the data can be found in table 5.2 where sum of squares (SS), mean square (MS), degrees of freedom (df), and F statistic for between and within group variations are shown. F value more than F critical rejects the null hypothesis for the corresponding significance level. For significance level $\alpha = 0.05$, the corresponding F critical value is 3.10. Since $F = 5.05$ is greater than F critical, the null hypothesis can be rejected meaning at least one group mean is different from the rest. Knowing that at least one population is different, further analysis is done on pairwise comparison of the treatment groups.

5.5.2 Pairwise Comparison & Tukey Process

Naturally, following the same line of analysis, the next step in pairwise comparison would be comparing all possible pairs of treatments, and finding the ones that are statistically significant.

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F Critical*
Between Groups	11.12	2	5.56	5.05	0.008	3.10
Within Groups	92.48	84	1.10			

Table 5.2: ANOVA Results.

However, doing so will increase the likelihood of observing rare events and rejecting the null hypotheses incorrectly, i.e., type I error [148]. There are a variety of methods to handle pairwise comparisons without increasing the chance of type I errors, among which Bonferroni correction and Tukey's method [149] are widely used with Tukey's method preferred slightly more in comparing all possible pairwise comparisons [1].

Tukey's method is based on calculating the value q from equation 5.1. Where μ_1 and μ_2 are the mean of gruop one and two, respectively, and SE is the standard error of the sum of the means.

$$q = \frac{|\mu_1 - \mu_2|}{SE} \tag{5.1}$$

Similar to ANOVA, if the q value for any pairwise comparison is greater than the $q_{critical}$, the null hypothesis for that comparison is rejected, i.e., the means of the two groups are different.

Comparison	*Mean Difference*	*SE*	*q*
Comfortable vs. Uncomfortable	0.75	0.19	3.89
Uncomfortable vs. Changing	0.75	0.19	3.89
Changing vs. Comfortable	0.00	0.19	0

Table 5.3: Tucky process data.

Table 5.3 shows the q values calculated for all the pairwise comparison. For $\alpha = 0.05$, 84 degrees of freedom, and three groups, $q_{critical}$ is 3.37. Therefore, the null hypotheses for comfortable vs. uncomfortable, and uncomfortable vs. changing is rejected; however, the experiment fails to reject the null hypothesis for changing vs. comfortable.

5.5.3 Session Order & Two-Way ANOVA

So far in our analysis, we have focused on finding out whether the between groups differences are statistically significant. In the next steps, we focus on quantifying the differences and possibly creating a mathematical model of user willingness that can be incorporated in RL algorithms. Throughout the experiment all external factors have been controlled for each scenario with the

[1]Engineering Statistics Handbook: https://itl.nist.gov/div898/handbook/prc/section4/prc473.htm

exception of the orders of the scenarios which were randomly assigned for each participant. Following the same line of analysis, a two-way ANOVA is performed on the data to test whether order of the treatment type creates statistically significant impact on the willingness ratings. Similar to one-way ANOVA, a two-way or two-factor ANOVA tests whether two different factors and their interactions create significant differences in the groups means. Specifically, two-way ANOVA tests the following sets of hypotheses:

- H_0^1: The means of observations grouped by the first factor are the same.
- H_a^1: At least one group mean is different from the rest when grouped by treatment.
- H_0^2: The means of observations grouped by the order are the same.
- H_a^2: At least one group mean is different from the rest when grouped by order.
- H_0^3: There is no interaction between treatment and order.
- H_a^3: Interaction exists between treatment and order.

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F Critical*
Treatment	10.17	2	5.08	6.03	0.003	3.12
Order	6.09	2	3.04	3.61	0.03	3.12
Interaction	3.60	4	0.90	1.06	0.37	2.49

Table 5.4: Tow-way ANOVA results.

Table 5.4 shows the results for the two-way ANOVA. The first null hypothesis is basically the one rejected by ANOVA test previously. The second null hypothesis is also rejected as the F statistic value is greater than F critical, meaning the order has impact on willingness ratings. Finally, the experiment fails to reject the last null hypothesis, i.e., the data does not support the claim that the interaction between order and treatment impacts the willingness. Finally, knowing what variables to include, a regression model can be fitted on the data to quantify the impact of each variable on the willingness ratings.

5.5.4 Regression Analysis

A linear regression is performed on the data to capture the first order trends using the variables identified previously. Tow-way ANOVA test showed that the treatment and order of the treatment are statically significant predictors of the user willingness for providing feedback. The interaction variable of treatment and order is not included as the two-way ANOVA failed to reject the third null hypothesis.

Order is an ordinal variable which can be directly included in the regression model. However, treatment variable is a categorical variable and has to be included using a the dummy variable method. In this method, two binary variables are introduced that can be either zero or one. Table 5.5 shows the two dummy variables TC and TU, and their value for each level of the treatment variable.

Treatment	TC	TU
Comfortable	1	0
Uncomfortable	0	1
Changing	0	0

Table 5.5: Dummy variables' values for the levels of treatment variable.

Equation 5.2 shows the regression model for willingness as a function of treatment dummy variables and order. The goal of regression analysis is to estimate the β coefficients along with their statistical significance.

$$W(TC, TU, Order) = \beta_{TC} * TC + \beta_{TU} * TU + \beta_{Order} * Order + Intercept \tag{5.2}$$

Table 5.6 shows the regression analysis results. The coefficients for order and TU have significant p-values as expected. TC dummy variable distinguishes between changing and comfortable levels for the treatment variable. The p-value for β_{TC} is not significant which is in line with the pairwise comparison of Tukey process.

Coefficients	*Estimate*	*SE*	*t value*	*P-value*
Intercept	5.08	0.33	15.26	$< 2e^{-16}$
β_{TC}	-0.02	0.26	-0.089	0.92
β_{TU}	-0.75	0.26	-2.84	0.0056
β_{Order}	-0.34	0.13	-2.56	0.01

Table 5.6: Regression analysis results.

5.6 Discussion

The most notable result is that the behavior of the algorithm in exploring/exploiting the design space impacts reported user willingness to interact with the adaptive system over long periods. In particular, desirability of the system state is a significant predictor of the interaction willingness. Theory from human factors and survey design suggests these factors would increase trust and reduce survey fatigue. However, the repeated preference elicitation task at the heart of the proposed adaptive systems based on RL algorithms are significantly different from the tasks previously studied. Results from the ANOVA tests and regression analysis indicate that for this particular office chair system, users were more willing to continue interacting when the chair was in a comfortable state. This may be due to perceived competence in the system. The system is "working" as intended and therefore should be trusted to use future information appropriately. Uncomfortable states may represent a "broken" system which could lower motivation for future interactions.

In contrast, system responsiveness was not important to user willingness under certain conditions. The experiment was unable to reject the hypothesis that a responsive system was better for user willingness than a static comfortable state. This result suggests that when the system is at a desirable state, the system does not need to be responsive in order to have continued user interaction. Given the tendency of RL algorithms to stay near desirable states, this finding suggests the current RL behavior may not be as detrimental to user willingness as hypothesized.

Another result corroborating prior work was the finding that the order of the experimental condition had a statistically significant negative effect on user willingness. As the null hypothesis was not rejected for the interaction terms, this suggests that user willingness to interact was reduced as the experiment continued and the subjects responded to more requests for feedback regardless of the treatment order and type. This is to be expected based on survey fatigue and human factors literature.

These results taken together have many implications for the design of future adaptive systems. Much like survey questionnaires, careful design of the elicitation mechanism could lead to significantly higher quality data. These findings suggest that RL algorithms used in this application could benefit from adaptive sampling methods to mitigate possible survey fatigue or damaged trust. The magnitude of the impact of the identified factors is likely due to the specific task and context. However, the approach used creates a road map for future designers to test these factors in their setting. For example, an adaptive hardware design team could perform a similar regression analysis and use the coefficient to determine the RL algorithm characteristics which maximize the probability of user interaction.

5.7 Conclusion & Future Work

This study presents results from an empirical pilot study investigating factors which impact user willingness to repeatedly interact with adaptive systems. A statistical analysis of reported willingness to interact under different experimental conditions suggests the following answers to the stated research questions.

1. Does user preference for the current system state impact user willingness to continue interacting with an adaptive system?

   Yes, in the case of the tested adaptive office chair, being in an uncomfortable system state had a large statistically significant negative effect on user willingness to continue interacting compared to the other two scenarios.

2. Does user perception of system responsiveness impact user willingness to continue interacting with an adaptive system?

   Under certain conditions, the system responsiveness did not impact user willingness to interact. In the case of the tested adaptive office chair, the null hypothesis was not rejected when comparing the static comfortable state to the responsive treatment.

In conclusion, designers of adaptive systems using RL algorithms should design systems to mitigate survey fatigue in order to ensure the required flow of feedback data. In doing so, they should consider using adaptive sampling techniques to mitigate negative effects on user willingness from undesirable states.

The study is limited by a number of factors. First, as a pilot study the number of participants was small. Although the experiment was able to adequately test many of the hypotheses, further study is needed to more precisely measure these effects. Second, the study was conducted in laboratory conditions which almost certainly affected the results. Users testing the comfort of their own office chairs over long periods of time may provide a more accurate picture of these phenomena. Finally, the frequency of feedback was quite short due to time constraints. It is uncertain whether the observed impacts would persist at different frequencies.

CHAPTER 6

Contributions, Implications, and Conclusions

6.1 Summary of Findings & Contributions

This dissertation is comprised of three studies proposing new frameworks for handling dynamic user preferences. Current design approaches do not respond quickly to changes in user preferences leading to lower user satisfaction over the lifetime of the product. Furthermore, current design frameworks like decision-based design (DBD) treat the usage context variables reactively. New frameworks are required to offer tools for designing products that have system level impact on the usage context, and changing the reactive role of designer to proactive with regard to usage context. Therefore, the presented studies focus on two sources of dynamicity in user preference: exogenous, and endogenous factors. First study proposes a framework for addressing exogenous factors impacting user preference, i.e., usage-context, and the other two studies focus on providing solutions for handling endogenous factors by proposing reinforcement learning based algorithm for design of "online" adaptive hardware systems.

Chapter 3 proposes a framework built on DBD and usage context-based design (UCBD) to assist designers with making products that have system level impact on the usage context through coupling mechanisms. Examples of this type of products are well explained in sustainability research on rebound effect. In this framework, an evolution function is introduced to the optimization model of UCBD. Evolution function predicts the system level impact of product attributes and updates the system level usage context variables. The updated variables are then iteratively fed back to a multidisciplinary feasible optimization architecture to find the optimal steady state of the product-user-environment. A vehicle mass design case study was used to demonstrate the value of the framework. Results from the case study suggest that the optimal vehicle mass can be significantly different depending on whether system-level impact of product attributes are incorporated in the design process or not. The system level impact of vehicle mass (product attribute) is modeled through the induced demand due to fuel economy elasticity (the evolution function). Fuel economy elasticity and induced demand are well explored mechanisms in literature, and using this example, the proposed framework offered a promising tool for supporting designers making decision which involve these types of mechanisms.

Chapter 4 fills the gap in literature on design for dynamic preference due to endogenous factors. Traditionally, designers develop products with the underlying assumption that the user preferences remain static during the development process or even the lifetime of the product. However, cases have been identified by prior works in which user preference is dynamic and preferred product attributes are changing as the user interacts with the product. This study proposes a theoretical framework based on reinforcement learning algorithms for designing adaptive physical products. These products constantly interact with the user through their lifetime and update their functionality to adapt to changes in user's preference. An illustrative case of design of a variable stiffness prosthetic ankle (VSPA) is presented to showcase the utility of the proposed framework. Results show that reinforcement learning algorithms can outperform traditional design methods under different scenarios and with no prior information. In particular, Upper Confidence Bound algorithms outperformed the rest of algorithms and design methods for the case study.

Chapter 5 builds upon the work done in chapter 4 on using reinforcement learning algorithms to design adaptive engineered systems. RL algorithms require a constant flow of preference feedback from the user to operate effectively. Therefore, it is important to identify the factors impacting user willingness to interact with the adaptive engineered systems. This study focuses on different exploration/exploitation behavior aspects of RL algorithms and their impact on user's willingness to provide prolonged feedback. Using an empirical study to answer the research questions, 29 subjects participated in a pilot study providing interaction willingness score for an adaptive office chair under different scenarios. The scenarios explored were 1) comfortable configuration, 2) uncomfortable configuration, and 3) changing configuration. Results show that desirability of the system state affects the user willingness to interact over long periods of time. Results also found no significant difference between responsive system state and an unresponsive desirable state. These findings help with design of better algorithms that maximize user satisfaction over time while minimizing burden of providing feedback (or equivalently maximizing willingness to interact with the adaptive engineered system.)

This work is an attempt to push engineering design methods toward online processes and make preference incorporation an ongoing process by creating adaptive hardware systems that continuously gather information on user preferences. As discussed in the background chapter, borrowing the concept of "online vs. offline" learning, traditional design methods such as decision-based design and design optimization lie generally at the offline end of of the spectrum of online-offline user preference incorporation methods. Supervised by human designers, these methods rely on gathering data on the target user periodically, and synthesizing the data to construct user preference functions. The constructed functions are then incorporated in the design process of the next generation of products. These three studies come together to address this gap in engineering design methods and push them toward "online" processes by proposing frameworks for design for dynamic user preferences, Fig. 6.1.

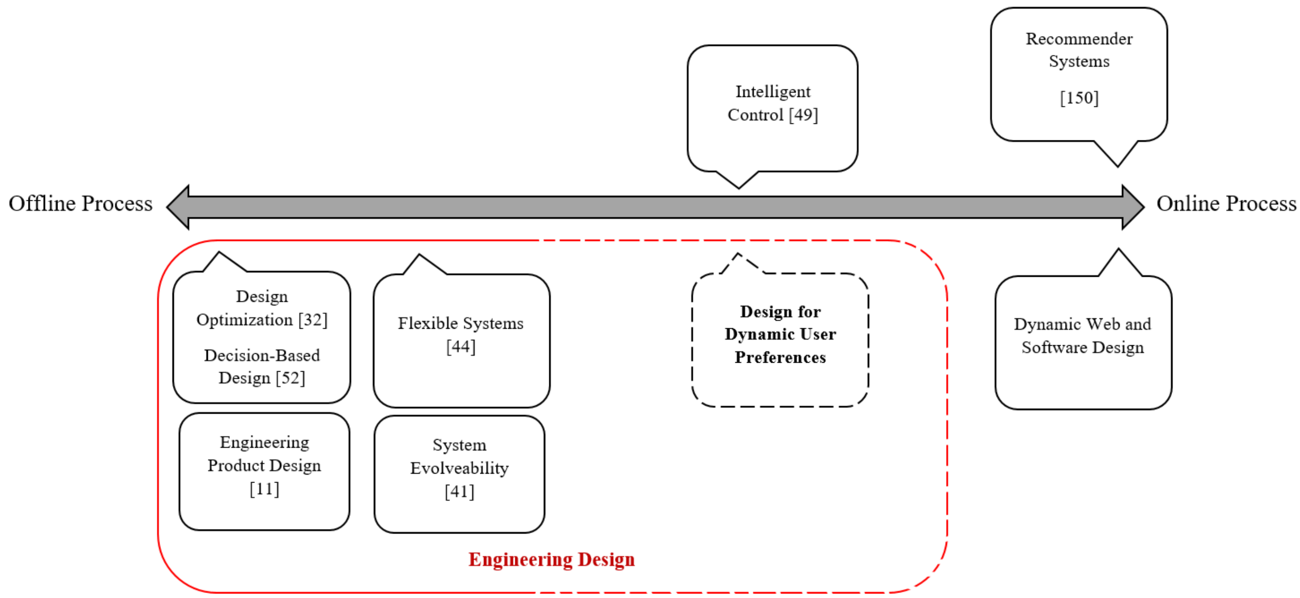

Figure 6.1: Incorporation of user preferences into product design.

6.2 Implications

There are many implications imagined for this work in engineering design practice. The framework proposed in the first study can easily be extended to model any system level impact of designer's decision on the environment. For example, using the same optimization architecture and design framework, sustainability goals can be incorporated in the design process as objective functions. As an example, rebound effect can be easily mitigated if the system level impact of the product is modeled as well in the design process using the proposed framework.

Results from second study demonstrate the utility of the framework for mass personalization purposes. Coming from software development, machine learning methods are widely used for mass personalizing software and websites. Using the same algorithms for dynamic preferences the framework can be easily extended to include mass personalizing adaptive hardware systems, i.e., the algorithms can react to changes in individual's preferences over time as well as the difference in preferences across a population of users.

Designing for the lifetime of a product means maximizing user preference and satisfaction over the entire product usage lifetime, as opposed to traditional design methods that mostly focus on maximizing preference at the time of purchase. This approach to product design ensures longer usage life which in turn helps with creating more sustainable products and solutions. In other words, having a product that adapts to changing user preferences reduces the need for switching to new product generations periodically, thus reducing the environmental footprint of product manufacturing.

Specific industries and business can benefit from the work presented in this dissertation. For example, devices used in healthcare rehabilitation often require constant adjusting and tuning as the recovery can be relatively fast and the state and preferences of the users may change rapidly throughout the recovery process. Taking advantage of the framework proposed for dynamic preferences, medical devices can adapt quickly to changes in the body and preferences of the patients as they interact with the hardware. Finally, this framework facilitates the transition from the traditional one-time sale business model to product as a service, subscription, and recurring revenue business models. These business models are well implemented and tested in software industry generating new types of value for both users and the companies.

6.3 Future Work

This dissertation provides examples of new approaches to design for dynamic preferences which has not been addressed by engineering design community previously. Although the results and examples are limited and preliminary, they open up many avenues for further investigation into design for dynamic user preference. For instance, in the case of design of a variable stiffness prosthetic ankle (VSPA), the design problem had only one variable, the ankle stiffness. Further work has to be done on cases with large number of design variables. In particular, new algorithms have to be developed for high dimensional design spaces that perform well while keeping the number of user interactions minimal. Moreover, in the examples explored, the algorithms only decided on the values of a predetermined set of attributes and not on what attributes to add and explore. Future work should be done on recommender algorithms [150] that not only pick the best value for each product attribute (design variable), but also recommend adding or removing attributes to/from the product.

In the third study, user willingness for providing feedback was explored under different scenarios. The focus of the study was on main RL algorithms' behaviors, namely high and low performance, and responsiveness. Future work has to be done on combinations of the main behavior and even more complex ones to measure the impact of algorithm exploration/exploitation behavior on user willingness for providing feedback under more complex and realistic scenarios. Moreover, although this study was able to capture the effect of usage duration on willingness, future work should explore how longer time scales may change the impact of algorithm behavior on user willingness for providing feedback. Finally, smart algorithms should be developed to measure the preference as well as user willingness for providing feedback (through response rate) and adapt their sampling method to keep the user engaged throughout the lifetime of the product.

6.4 Conclusion

This dissertation contributes to the engineering design field by extending the current design methods and bringing new ones from machine learning community to address design for dynamic preferences. The source of dynamicity can either come from individual's change in preference for the same product attributes under the same environmental conditions, or may come from changes imposed to the user's preferences due to changes in contextual factors. The importance of addressing dynamic user preferences is shown by presenting different case studies and comparing them to solutions provided by current engineering design practices. Moreover, the examples provided in this study open up new avenues for addressing mass personalization as well as design for sustainability. This dissertation will create new opportunities for exploiting the abundance of data and data gathering tools in order to improve product design. Starting from the usage context, a framework is proposed for designing the product and the usage context simultaneously. A framework is then proposed to change the role of designer from designing the actual product to designing the policies and algorithms used in product design. The goal is to enable the product to cycle through the design stages starting from preference elicitation and ending in the final product without designer's supervision. In conclusion, this research provides a step toward pushing engineering design methods closer to online processes by bringing new design tools from other disciplines to engineering design.

APPENDIX A

Traffic Simulation Platform

The simulation platform consists of four modules namely, map optimization module, demand module, agent behavior, and path finding module. In this appendix, each module will be explained in detail.

A.1 Location/Map and Road Network

As an open-source platform, Open Street Map (OSM)[1] gives free access to an editable map of the world that is built and maintained by volunteers. Due to the platform's accessibility along with the availability of third-party Python packages, OSM was used to import maps for the map optimization module.

The maps can be imported in a variety of ways such as importing by defining the four coordinates of the box surrounding the desired area, by address, and by defining the center coordinate and radius of the area of interest. Figure A.1 shows a map of the part of southeast Michigan that will be used for the simulations.

The graph generated by the OSMx Python package gives a complete graph of the map of the area which had to be further refined for our purposes [151]. Figure A.2 shows the initial graph imported of the area shown in figure A.1.

A set of refinements is done on the initial graph to make it usable for the model. The first step is to limit the nodes and links to the level of details deemed sufficient for simulations. Table A.1 shows the road levels and their definitions included in the map. For the road network, the motorway, trunk, primary, secondary, tertiary, residential and unclassified roads are included.

After choosing the right level of details for the road system, we have to ensure that the resulting graph stays a connected graph, where any trip between any two nodes remains possible. Therefore, a code is run to select a connected subgraph of the road system. Figure A.3 shows the connected subgraph of the graph shown in figure A.2 for the road network

[1]https://www.openstreetmap.org/

Road Type	*Definition*
Motorway	A restricted access major divided highway, normally with 2 or more running lanes plus emergency hard shoulder. Equivalent to a freeway, Autobahn, etc.
Trunk	The most important roads in a country's road system that are not motorways. (Not necessarily be a divided highway.)
Primary	The next most important roads in a country's system. (Often link larger towns.)
Secondary	The next most important roads in a country's system. (Often link towns.)
Tertiary	The next most important roads in a country's system. (Often link smaller towns and villages)
Unclassified	The least important through roads in a country's system – i.e. minor roads of a lower classification than tertiary, but which serve a purpose other than access to properties. (Often link villages and hamlets.)
Residential	Roads which serve as an access to housing, without function of connecting settlements. Often lined with housing.

Table A.1: Road levels and definitions.

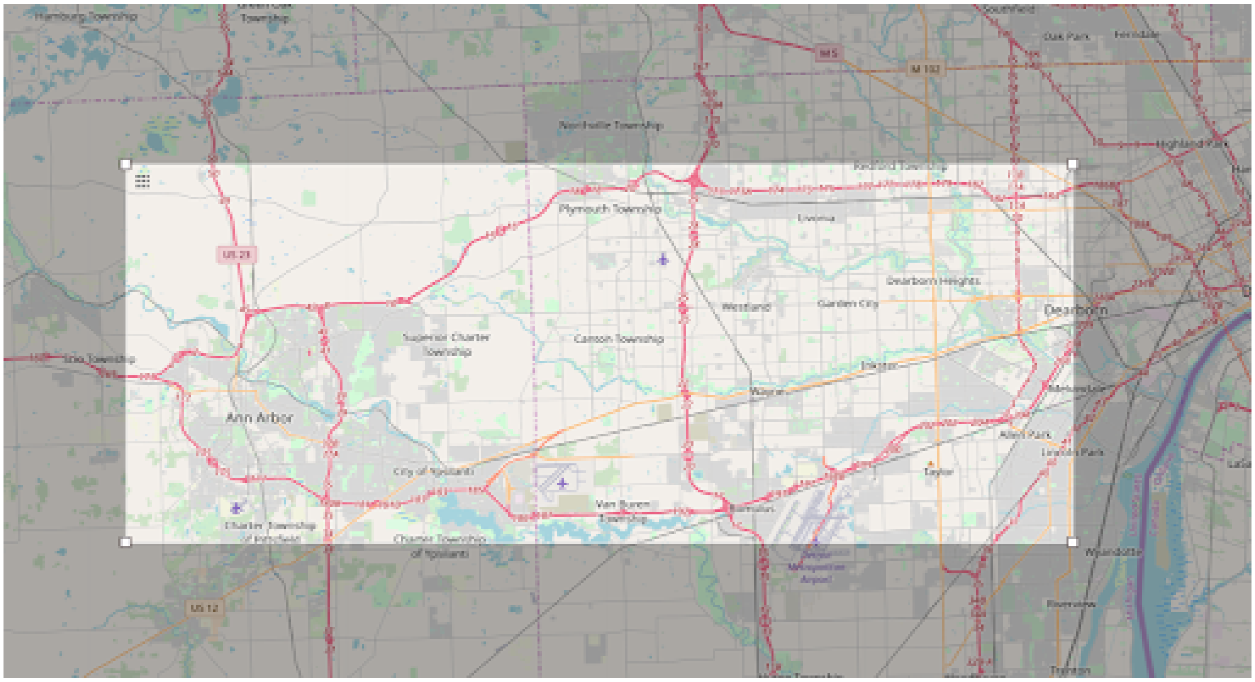

Figure A.1: Simulation runs input map.

As it can be seen from figure A.3, there are many nodes that are not intersections and can be removed from the graph without significant loss of information. This removal is particularly important for the pathfinding module as smaller numbers of nodes and edges lower the computation cost of calculating an optimal path between two points. Figures A.4 shows the detailed map after node simplification.

A.2 Demand

Using technical reports from the Southeast Michigan Council of Governments (SEMCOG), we create different types of travel demands. A trip is defined as a movement from one location to another without having a stop. Vehicle trips per person is the number of vehicle travels per person in a day. By modeling the demand based off of vehicle trips per person we model the demand with any number of passengers implicitly. Figure A.5 shows the vehicle trips per person for the counties in the south east Michigan counties [2].

Demands are generated by assigning 1) Starting time 2) a starting node 3) demand type 4) travel length (miles) 5) destination node. Prior to generating the demand, each node is assigned a population. Using US census data, we first find the population of each census tract, figure A.6. Census tracts are the highest resolution of data regions in the census data. Once the population

[2] SEMCOG, Travel Characteristics: Technical Report, Oct 2016

Figure A.2: Initial road system graph of area shown in figure A.1.

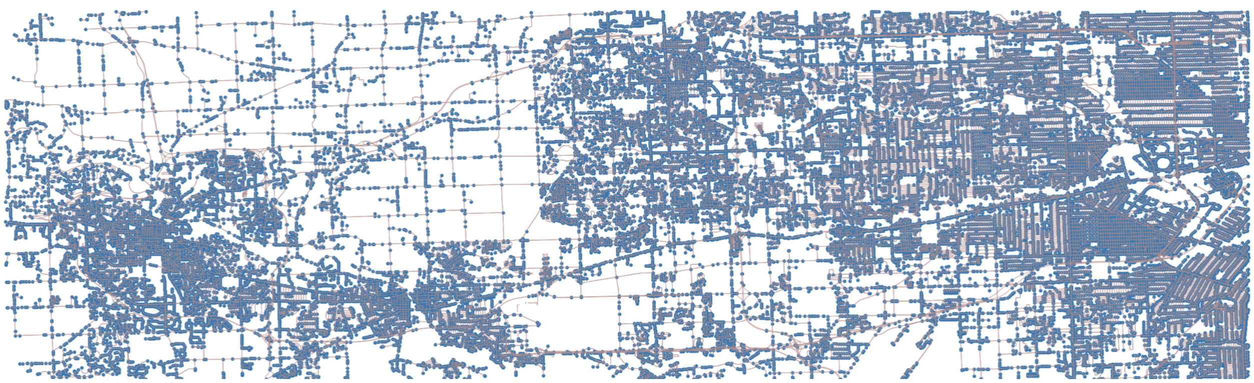

Figure A.3: Connected subgraph of the road system of figure A.2.

of each tract is found, the tract population is then distributed to all the nodes within the tracts excluding the highway nodes, figure A.7.

To have a more spreaded population distribution, every 200 meters extra nodes are added to the road network. Figure A.8 shows the node distribution for Ann Arbor.

Once the node populations are assigned, vehicle trip per node is found by multiplying each node's population by its vehicle trip per person. Figure A.8, shows nodes' vehicle trip per person found from figure A.5.

Next step is to generate demand starting times. Using figure A.10[3], at each second during the simulation, the number of trips starting from each node is calculated using Poisson distribution. The simulation starts from 2 pm, and ends at 2 am; however, the demands are generated until midnight and amount to 898222 trips. We have only modeled the demands starting time distribution

[3]mdot: MI Travel Counts III, Travel Characteristics, Technical Report, Sep 2016

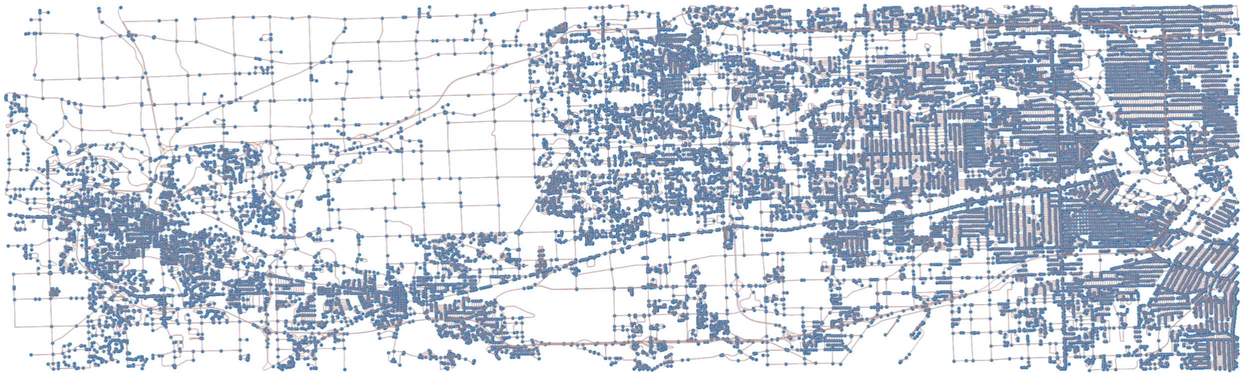

Figure A.4: Road network after simplification.

Table 1-1. **Number of households, people, and weekday trips by region (weighted)**

Region	Households (wtd)	Persons (wtd)	Trips per person	Vehicle trips per person	Trips per household	Vehicle trips per household
Washtenaw Area (WATS)	136,471	332,614	4.0	2.3	9.6	5.4
Eastern Wayne	287,299	774,076	3.8	1.7	10.1	4.6
Western Wayne	380,256	994,861	3.6	2.4	9.5	6.3
Oakland	489,635	1,208,898	3.8	2.5	9.4	6.2
Macomb	334,509	841,088	3.7	2.4	9.5	6.2
Monroe	58,328	149,499	3.7	2.4	9.9	6.6
St. Clair	64,182	159,325	3.8	2.5	9.7	6.2
Livingston	68,439	182,607	3.7	2.5	9.5	6.5
Total	1,819,119	4,642,968	3.8	2.3	9.6	5.9

Figure A.5: Number of households, people, and weekday trips by SE Michigan counties.

based on figure A.10. Figure A.11[4] shows the types of demand during the morning and afternoon pick hours. Using the data on the afternoon pick hours, we modeled four main types of demand, namely, commute, shop/errands, visit/social, and pick-up/drop off.

Demands are randomly assigned a type from the four types mentioned based on the afternoon peak hour trip purpose distribution. The trip types are round trips except for commute trips which are one-way trips. The round trips are basically two demands with a time gap between the end of the first trip and start of the second trip. The time gap is calculated randomly using a normal distribution. The time gap for shopping is calculated based on the average shopping time of 43 minutes[5]. Since no data was found on the time gaps of other types of trips, we assumed an average of 30 seconds and 3 hours for pick-up/drop off and visit/socialize, respectively.

[4]SEMCOG, Travel Characteristics: Technical Report, Oct 2016

[5]https://www.ers.usda.gov/amber-waves/2020/april/more-americans-spend-more-time-in-food-related-activities-than-a-decade-ago/

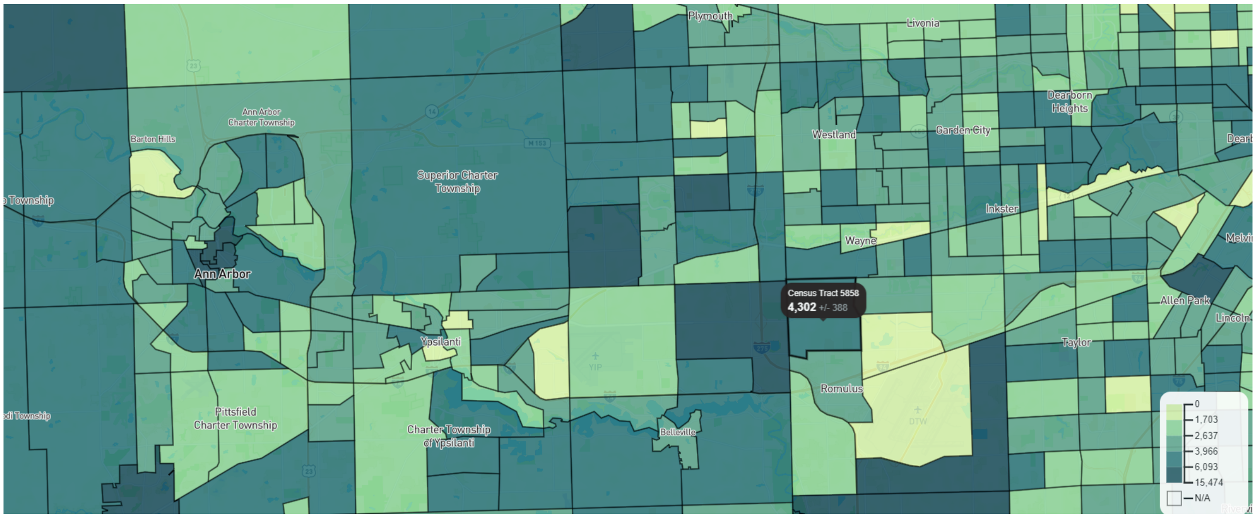

Figure A.6: Census tracts total population.

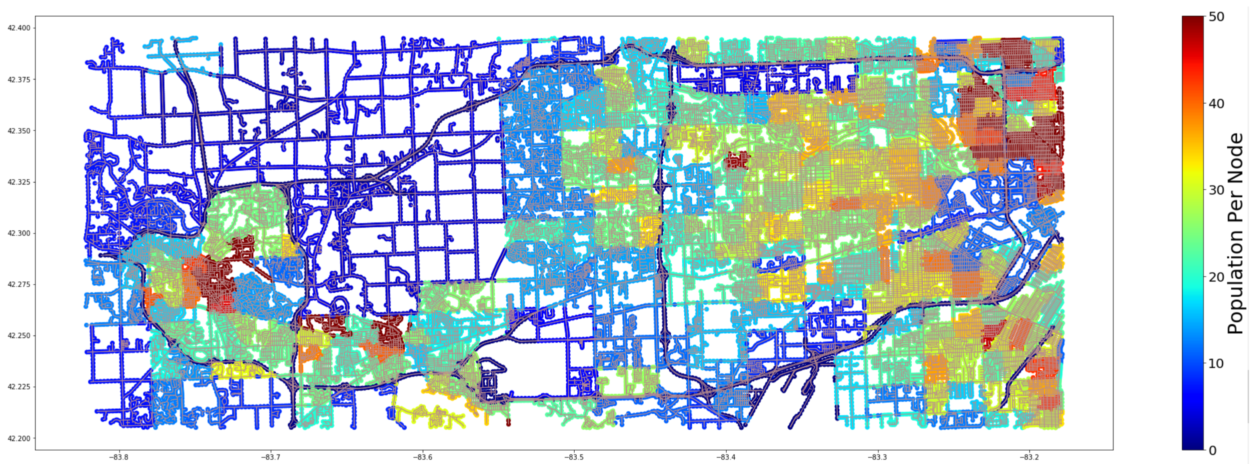

Figure A.7: Census tracts total population.

In the final step, the destination nodes are assigned for each demand. Depending on the demand type, first the algorithm randomly chooses a trip length (miles) following a normal distribution. Once the trip length is picked, all the nodes within a tolerance length of the chosen length are found using Dijkstra's algorithm with road lengths as edge weights. The destination node is then picked randomly from the set of the nodes found in the previous step. The tolerance is set at 200 meters. The average trip length for commuting is 13.2 miles[6] and for shopping is 4 miles[7] . No data was found on the average pickup and visit travel lengths, and we assumed a 10 mile average for both. The destination nodes are further restricted by trip types. For example, shopping and

[6]SEMCOG, Travel Characteristics: Technical Report, Oct 2016

[7]https://usa.streetsblog.org/2015/04/10/5-things-the-usda-learned-from-its-first-national-survey-of-food-access/

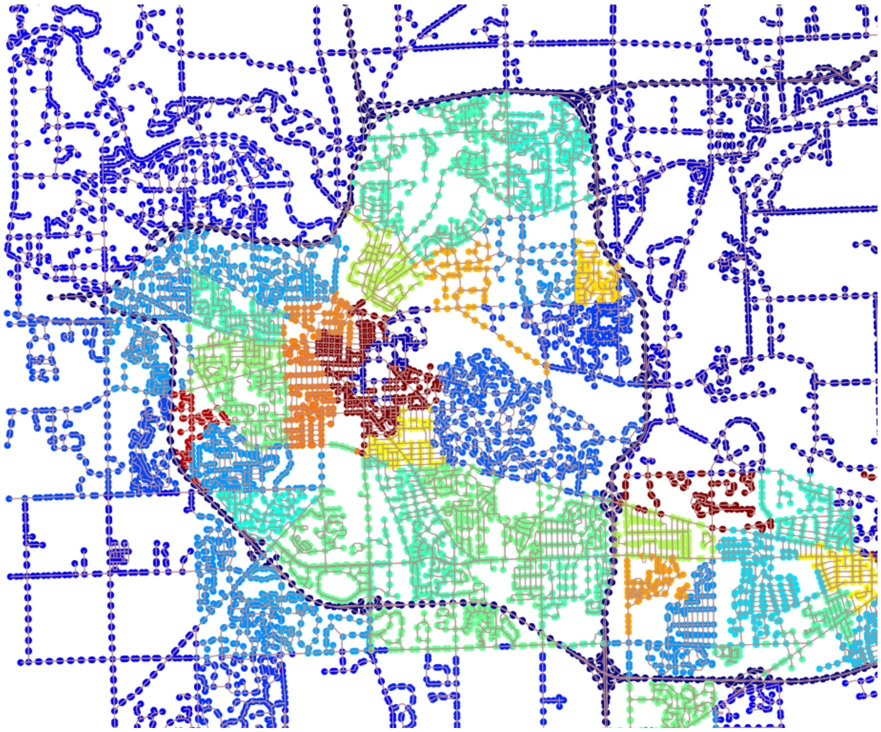

Figure A.8: Node distribution for Ann Arbor.

work locations cannot be in the residential area. Figure A.12, shows the residential only nodes and figure A.13 shows the acceptable nodes for shopping and work locations.

A.3 Agent Behavior

After exploring transportation modeling literature and experimenting with different car following models a set of simple rules were defined for each agent (driver) to behave based on. The simulation approach falls in the large framework of Agent Based Modelling (ABM) where a population of agents interact with each other based on a set of rules.

The set of agent behaviors is as follows:

1. Agents try to reach the speed limit.

2. Agents maintain 1 car length distance to the next car for every 10 mph of speed.

3. The acceleration/deceleration rates are fixed.

4. To prevent chaotic and unstable traffic flow, agents are not allowed to overtake the next car; however, they move to the lane with the fewest number of cars every 200 meters.

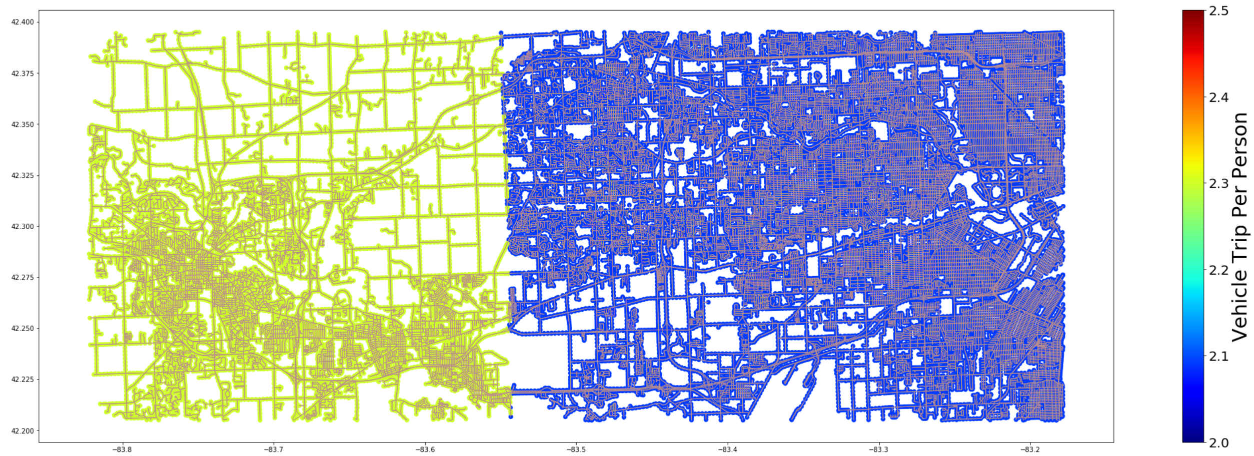

Figure A.9: Vehicle trip per person for each node in the simulation region.

To further decrease the run time of the model, constant acceleration motion formulas are used instead of general motion differential equations. Under the assumption of constant acceleration, this simplification enables us to choose relatively large time steps with minimal compromise in simulation accuracy. The final simulation time step is set at 0.25 second. Traffic lights are not modeled since without knowing the exact timing of each traffic light in the real world, the complexity added to the simulation may not be justifiable. Figure A.14 shows the simulated agents at a single intersection. Red agents are braking, and black and green ones are cruising and accelerating, respectively.

A.4 Path Finding

To find the optimal path between any two points, Dijkstra's algorithm is used [152]. Depending on the edge weights, Dijkstra's algorithm can give the both shortest path and fastest path. Although it gives the global minimum, Dijkstra's algorithm is computationally expensive as finding the shortest path between the source node and the destination node gives the shortest path between the source and all the other nodes as well.

Figure A.15 shows the fastest path between two random nodes using Dijkstra's algorithms. The edge weights are travel time calculated based on speed limits. Figure A.16 shows the fastest path for the two same nodes on Google Maps.

The weights for the graph used to calculate the fastest path change throughout the simulation as the traffic flow rate changes for each edge. Therefore, the edge weights used for Dijkstra's algorithm have to be updated as well. The travel time for each edge (road) is updated every time a car leaves the road segment. An exponential moving average updates each edge's weight weighing

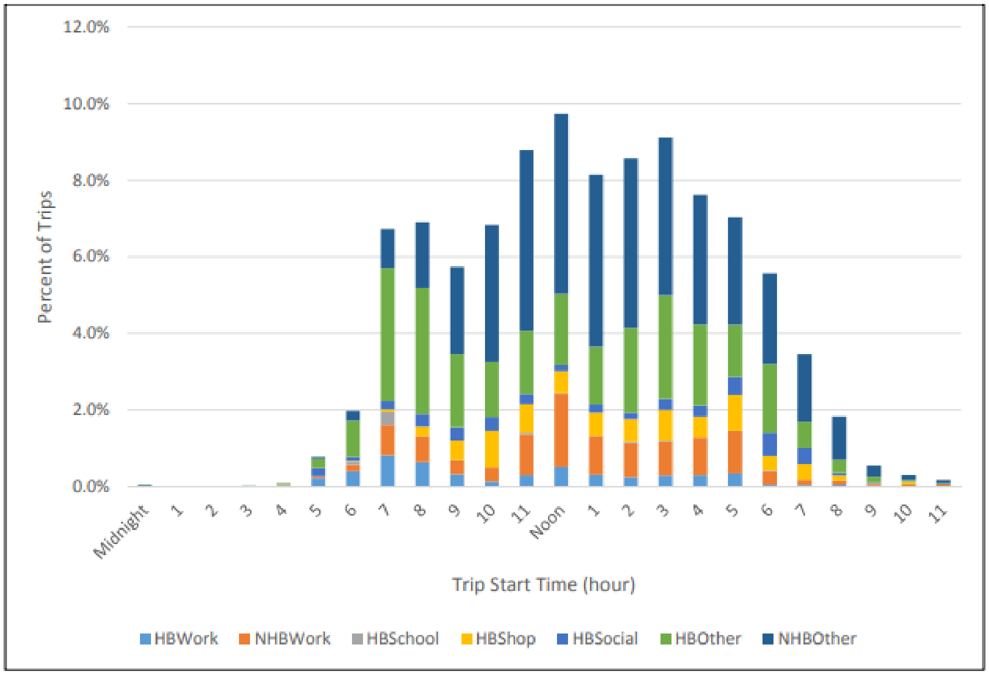

Figure A.10: Person trips by start time.

recent travel times more. Every 5 minutes the graph weights are updated by the moving average of cars' travel times for each edge.

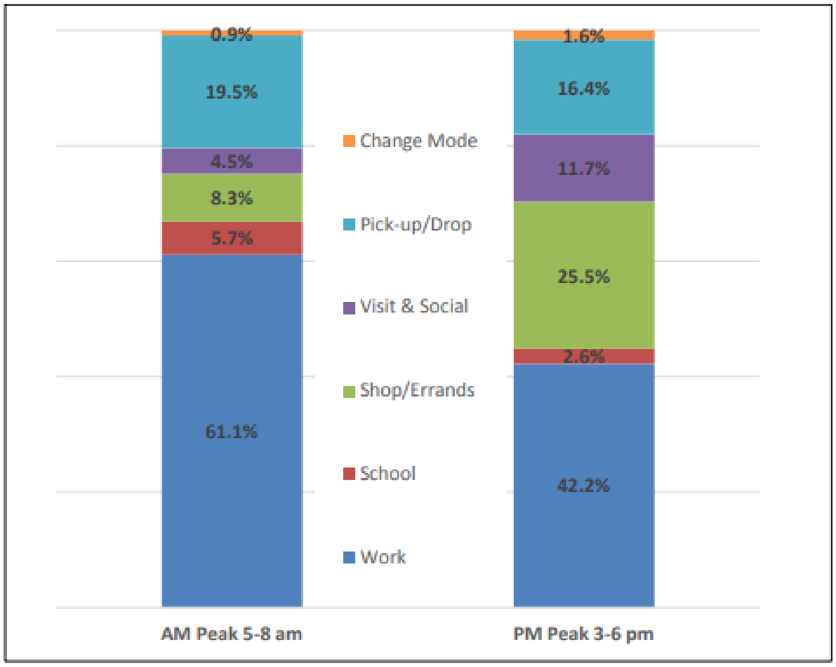

Figure A.11: Percent of vehicle trips by purpose.

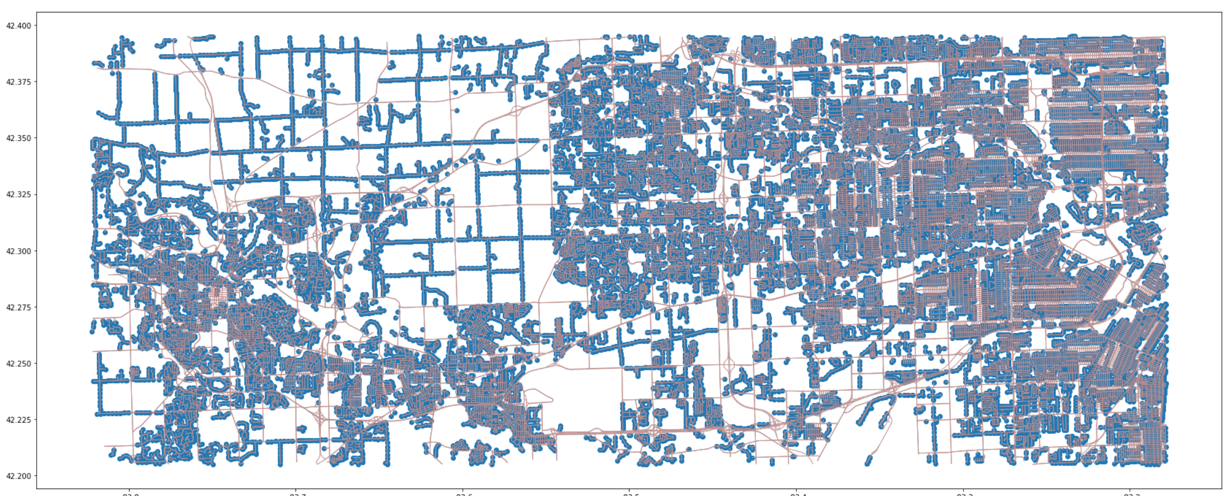

Figure A.12: Residential only nodes.

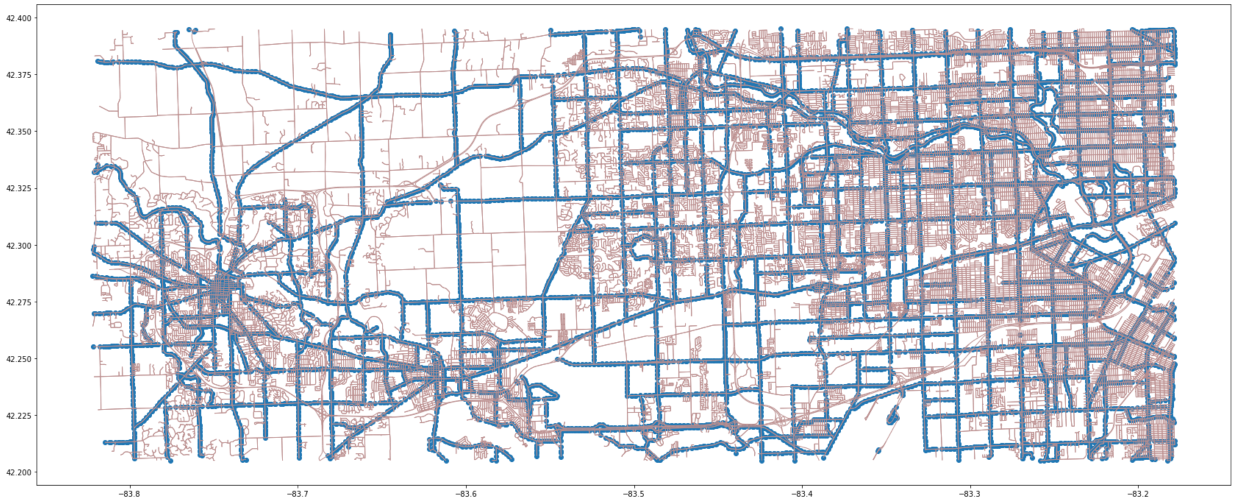

Figure A.13: Acceptable nodes for shopping and work locations.

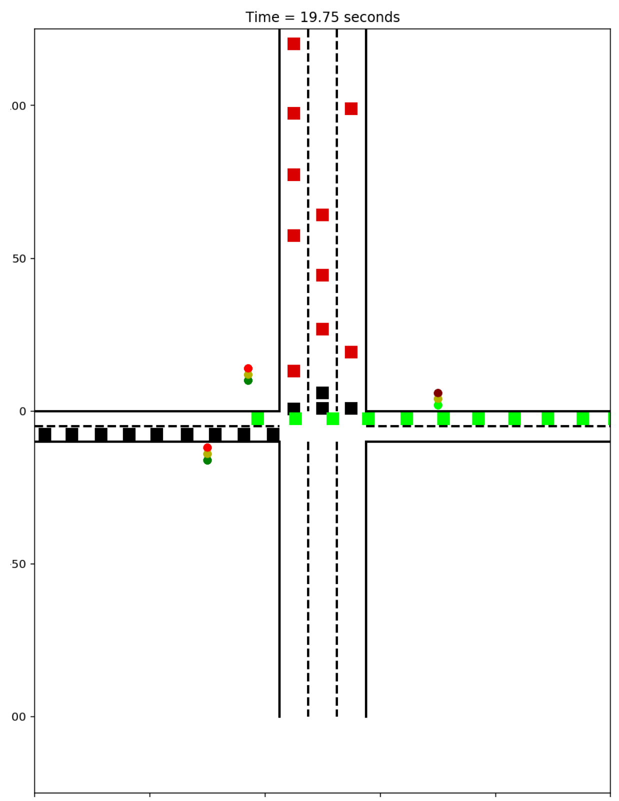

Figure A.14: Agent based modeling of one intersection.

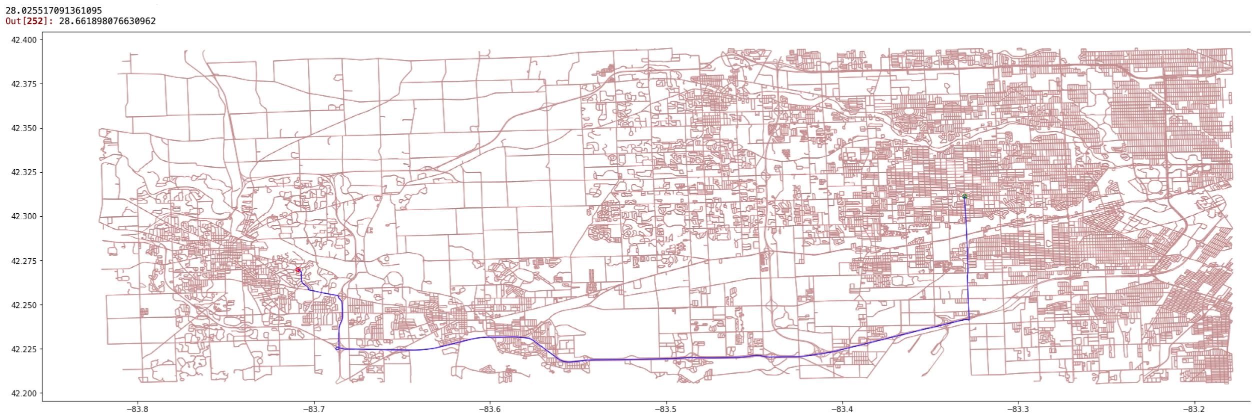

Figure A.15: Fastest path calculated by the path finding algorithm between two randomly chosen nodes.

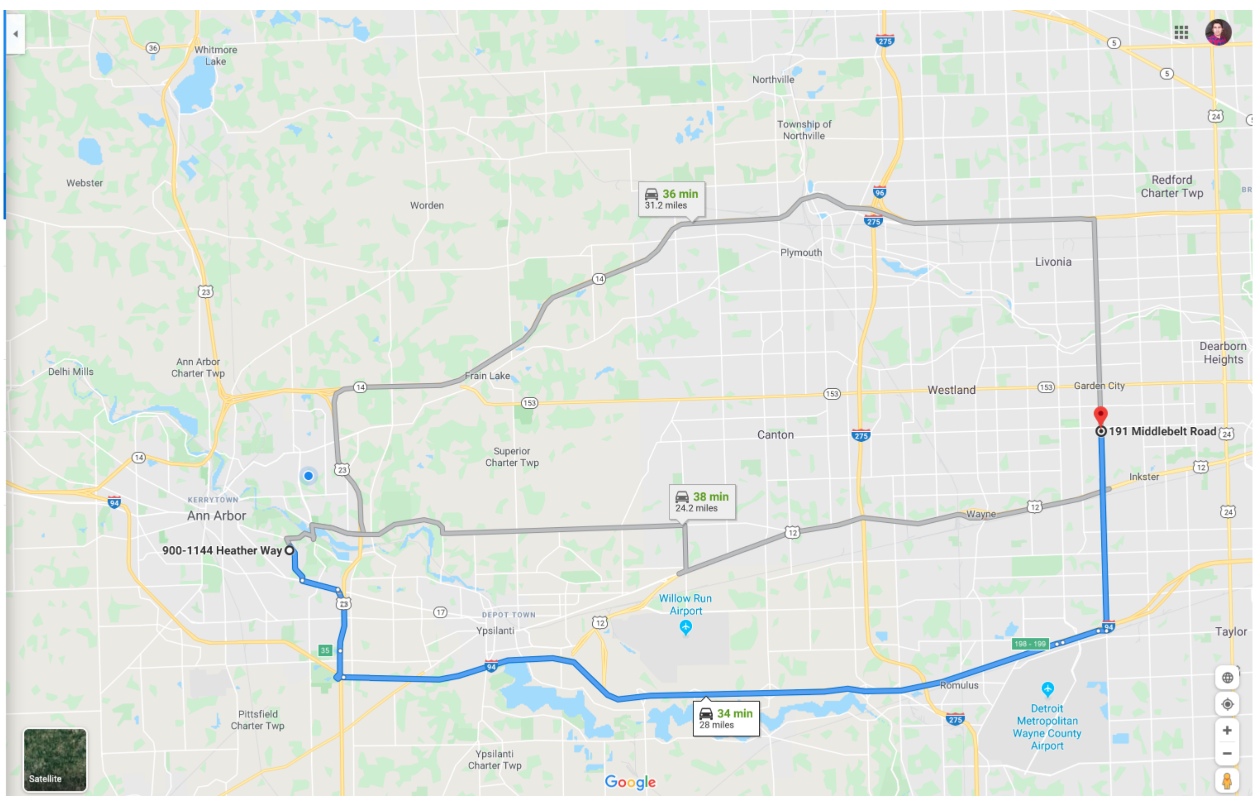

Figure A.16: Same path as fig A.15 found on Google maps.

BIBLIOGRAPHY

[1] Amit Sheth. Internet of things to smart iot through semantic, cognitive, and perceptual computing. *IEEE Intelligent Systems*, 31(2):108–112, 2016.

[2] Guy Harrison. Sharding, amazon, and the birth of nosql. In *Next Generation Databases*, pages 39–51. Springer, 2015.

[3] Dokyun Lee, Kartik Hosanagar, and Harikesh S Nair. Advertising content and consumer engagement on social media: Evidence from facebook. *Management Science*, 64(11):5105–5131, 2018.

[4] Philip J Trocchia and Swinder Janda. An investigation of product purchase and subsequent non-consumption. *Journal of Consumer Marketing*, 2002.

[5] Jacqueline R Rifkin and Jonah Berger. How nonconsumption can turn ordinary items into perceived treasures. *Journal of the Association for Consumer Research*, 6(3):000–000, 2021.

[6] Karl T Ulrich. *Product design and development*. Tata McGraw-Hill Education, 2003.

[7] John Broome. *Can a Humean be moderate?* University, 1990.

[8] George A Hazelrigg. A framework for decision-based engineering design. 1998.

[9] Beth E Allen. Toolkit for decision-based design theory. *Journal of Engineering Valuation and Cost Analysis*, 3(2):85–105, 2000.

[10] W Chen, KE Lewis, and L Schmidt. Decision-based design: an emerging design perspective. *Engineering Valuation & Cost Analysis, Special Edition on Decision Based Design: Status & Promise*, 3(1):57–66, 2000.

[11] Steven Eppinger and Karl Ulrich. *Product design and development*. McGraw-Hill Higher Education, 2015.

[12] Wei Chen, Carolyn Conner, and Bernard Yannou. User needs and preferences in engineering design. *Journal of Mechanical Design*, 137(7):070301, 2015.

[13] George A Hazelrigg. An axiomatic framework for engineering design. 1999.

[14] Alexander Brem, Volker Bilgram, and Adele Gutstein. Involving lead users in innovation: A structured summary of research on the lead user method. *International Journal of Innovation and Technology Management*, 15(03):1850022, 2018.

[15] Douglas Schuler and Aki Namioka. *Participatory design: Principles and practices*. CRC Press, 1993.

[16] Tomasz Miaskiewicz and Kenneth A Kozar. Personas and user-centered design: How can personas benefit product design processes? *Design studies*, 32(5):417–430, 2011.

[17] Stephen McCarthy, Paidi O'Raghallaigh, Simon Woodworth, Yoke Lin Lim, Louise C Kenny, and Frédéric Adam. An integrated patient journey mapping tool for embedding quality in healthcare service reform. *Journal of Decision Systems*, 25(sup1):354–368, 2016.

[18] Lai-Kow Chan and Ming-Lu Wu. Quality function deployment: A literature review. *European journal of operational research*, 143(3):463–497, 2002.

[19] John R Hauser, Don Clausing, et al. The house of quality. 1988.

[20] Madhan Shridhar Phadke. *Quality engineering using robust design*. Prentice Hall PTR, 1995.

[21] William Y Fowlkes and Clyde M Creveling. *Engineering methods for robust product design: using Taguchi methods in technology and product development*. Addison-Wesley, 1995.

[22] Louis L Thurstone. A law of comparative judgment. *Psychological review*, 34(4):273, 1927.

[23] R Duncan Luce. *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2012.

[24] Paul Edgar Green. New way to measure consumers' judgements. *Harv. Bus. Rev.*, 53:107–117, 1975.

[25] Paul E Green and Venkat Srinivasan. Conjoint analysis in marketing: new developments with implications for research and practice. *Journal of marketing*, 54(4):3–19, 1990.

[26] Paul E Green and Venkatachary Srinivasan. Conjoint analysis in consumer research: issues and outlook. *Journal of consumer research*, 5(2):103–123, 1978.

[27] Alex Burnap, Yanxin Pan, Ye Liu, Yi Ren, Honglak Lee, Richard Gonzalez, and Panos Y Papalambros. Improving design preference prediction accuracy using feature learning. *Journal of Mechanical Design*, 138(7):071404, 2016.

[28] Jarod Kelly, Panos Y Papalambros, et al. Use of shape preference information in product design. In *DS 42: Proceedings of ICED 2007, the 16th International Conference on Engineering Design, Paris, France, 28.-31.07. 2007*, pages 803–804, 2007.

[29] Seth Orsborn, Jonathan Cagan, and Peter Boatwright. Quantifying aesthetic form preference in a utility function. *Journal of Mechanical Design*, 131(6):061001, 2009.

[30] Kosa Goucher-Lambert and Jonathan Cagan. The impact of sustainability on consumer preference judgments of product attributes. *Journal of Mechanical Design*, 137(8):081401, 2015.

[31] David Van Horn, Andrew Olewnik, and Kemper Lewis. Design analytics: capturing, understanding, and meeting customer needs using big data. In *ASME 2012 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, pages 863–875. American Society of Mechanical Engineers Digital Collection, 2012.

[32] Panos Y Papalambros and Douglass J Wilde. *Principles of optimal design: modeling and computation*. Cambridge university press, 2000.

[33] Wei Chen, Christopher Hoyle, and Henk Jan Wassenaar. *Decision-based design: Integrating consumer preferences into engineering design*. Springer Science & Business Media, 2012.

[34] Henk Jan Wassenaar and Wei Chen. An approach to decision-based design with discrete choice analysis for demand modeling. *J. Mech. Des.*, 125(3):490–497, 2003.

[35] Jaroslaw Sobieszczanski-Sobieski. Multidisciplinary design optimization: an emerging new engineering discipline. In *Advances in structural optimization*, pages 483–496. Springer, 1995.

[36] Noah Williams, Shapour Azarm, and PK Kannan. Engineering product design optimization for retail channel acceptance. *Journal of Mechanical Design*, 130(6), 2008.

[37] Ferruh Erdogdu. *Optimization in food engineering*. CRC Press, 2008.

[38] Erin F MacDonald, Richard Gonzalez, and Panos Y Papalambros. Preference inconsistency in multidisciplinary design decision making. 2009.

[39] Morris A Cohen, Jehoshua Eliasberg, and Teck-Hua Ho. New product development: The performance and time-to-market tradeoff. *Management Science*, 42(2):173–186, 1996.

[40] David Van Horn, Andrew Olewnik, and Kemper Lewis. Design analytics: capturing, understanding, and meeting customer needs using big data. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 45066, pages 863–875. American Society of Mechanical Engineers, 2012.

[41] Morgan WP Tackett, Christopher A Mattson, and Scott M Ferguson. A model for quantifying system evolvability based on excess and capacity. *Journal of Mechanical Design*, 136(5):051002, 2014.

[42] Jeffrey D Allen, Christopher A Mattson, Kendall S Thacker, and Scott M Ferguson. Design for excess capability to handle uncertain product requirements in a developing world setting. *Research in Engineering Design*, 28(4):511–527, 2017.

[43] Scott Ferguson, Afreen Siddiqi, Kemper Lewis, and Olivier L de Weck. Flexible and reconfigurable systems: Nomenclature and review. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 48078, pages 249–263, 2007.

[44] Andrew Olewnik, Trevor Brauen, Scott Ferguson, and Kemper Lewis. A framework for flexible systems and its implementation in multiattribute decision making. *J. Mech. Des.*, 126(3):412–419, 2004.

[45] John K Gershenson, GJ Prasad, and Yan Zhang. Product modularity: definitions and benefits. *Journal of Engineering design*, 14(3):295–313, 2003.

[46] Waldemar Karwowski. Ergonomics and human factors: the paradigms for science, engineering, design, technology and management of human-compatible systems. *Ergonomics*, 48(5):436–463, 2005.

[47] Ryan Dew and Asim Ansari. Bayesian nonparametric customer base analysis with model-based visualizations. *Marketing Science*, 37(2):216–235, 2018.

[48] Peter Barna, Geert-Jan Houben, and Flavius Frasincar. Specification of adaptive behavior using a general-purpose design methodology for dynamic web applications. In *International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, pages 283–286. Springer, 2004.

[49] Duško Katić and Miomir Vukobratović. Survey of intelligent control techniques for humanoid robots. *Journal of Intelligent and Robotic Systems*, 37(2):117–141, 2003.

[50] Mohamed A Mabrok, Hassan K Mohamed, Abdel-Haleem Abdel-Aty, and Ahmed S Alzahrani. Human models in human-in-the-loop control systems. *Journal of Intelligent & Fuzzy Systems*, 38(3):2611–2622, 2020.

[51] Shai Ben-David, Eyal Kushilevitz, and Yishay Mansour. Online learning versus offline learning. *Machine Learning*, 29(1):45–63, 1997.

[52] Wei Chen, Christopher Hoyle, and Henk Jan Wassenaar. *Decision-based design: Integrating consumer preferences into engineering design*. Springer Science & Business Media, 2012.

[53] Vithala R Rao et al. *Applied conjoint analysis*. Springer, 2014.

[54] Mojtaba Arezoomand, Elliott Rouse, and Jesse Austin-Breneman. Theoretical framework for design for dynamic user preferences. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 84003, page V11AT11A016. American Society of Mechanical Engineers, 2020.

[55] Hui Li and Shapour Azarm. Product design selection under uncertainty and with competitive advantage. *J. Mech. Des.*, 122(4):411–418, 2000.

[56] Suzanne Chou, Mojtaba Arezoomand, Marianna J Coulentianos, Kowit Nambunmee, Richard Neitzel, Achyuta Adhvaryu, and Jesse Austin-Breneman. The stakeholder agreement metric: Quantifying preference agreement between product stakeholders. *Journal of Mechanical Design*, 143(3), 2021.

[57] Jeremy J Michalek, Fred M Feinberg, and Panos Y Papalambros. Linking marketing and engineering product design decisions via analytical target cascading. *Journal of product innovation management*, 22(1):42–62, 2005.

[58] Lin He, Wei Chen, Christopher Hoyle, and Bernard Yannou. Choice modeling for usage context-based design. *Journal of Mechanical Design*, 134(3), 2012.

[59] Russell W Belk. An exploratory assessment of situational effects in buyer behavior. *Journal of marketing research*, 11(2):156–163, 1974.

[60] Peter R Dickson. Person-situation: Segmentation's missing link. *Journal of marketing*, 46(4):56–64, 1982.

[61] Jaime Romero De La Fuente and María Jesús Yagüe Guillén. Identifying the influence of product design and usage situation on consumer choice. *International Journal of Market Research*, 47(6):667–686, 2005.

[62] Matthew G Green, Julie S Linsey, Carolyn C Seepersad, Kristin L Wood, and Dan J Jensen. Frontier design: a product usage context method. In *ASME 2006 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, pages 99–113. American Society of Mechanical Engineers Digital Collection, 2006.

[63] Lin He, Mingxian Wang, Wei Chen, and Guenter Conzelmann. Incorporating social impact on new product adoption in choice modeling: A case study in green vehicles. *Transportation Research Part D: Transport and Environment*, 32:421–434, 2014.

[64] Mingxian Wang, Wei Chen, Yun Huang, Noshir S Contractor, and Yan Fu. Modeling customer preferences using multidimensional network analysis in engineering design. *Design Science*, 2, 2016.

[65] Michael J Grubb. Communication energy efficiency and economic fallacies. *Energy Policy*, 18(8):783–785, 1990.

[66] Zhaohua Wang, Bai Han, and Milin Lu. Measurement of energy rebound effect in households: evidence from residential electricity consumption in beijing, china. *Renewable and Sustainable Energy Reviews*, 58:852–861, 2016.

[67] Kenneth A Small and Kurt Van Dender. The effect of improved fuel economy on vehicle miles traveled: estimating the rebound effect using us state data, 1966-2001. 2005.

[68] Joshua Linn. The rebound effect for passenger vehicles. *The Energy Journal*, 37(2), 2016.

[69] Murat Hakki Eres, Marco Bertoni, Mario Kossmann, and James Scanlan. Mapping customer needs to engineering characteristics: an aerospace perspective for conceptual design. *Journal of Engineering Design*, 25(1-3):64–87, 2014.

[70] Sergio A Brambila-Macias, Tomohiko Sakao, and Christian Kowalkowski. Bridging the gap between engineering design and marketing: insights for research and practice in product/service system design. *Design Science*, 4, 2018.

[71] Namwoo Kang, Fred M Feinberg, and Panos Y Papalambros. Integrated decision making in electric vehicle and charging station location network design. *Journal of Mechanical Design*, 137(6), 2015.

[72] R Aydin, Chun Kit Kwong, and Ping Ji. Coordination of the closed-loop supply chain for product line design with consideration of remanufactured products. *Journal of Cleaner Production*, 114:286–298, 2016.

[73] Jeremy J Michalek. Design for market systems: Integrating social, economic, and physical sciences to engineer product success. *Mechanical Engineering-CIME*, 130(11):32–37, 2008.

[74] Bart D Frischknecht, Katie Whitefoot, and Panos Y Papalambros. On the suitability of econometric demand models in design for market systems. *Journal of Mechanical Design*, 132(12), 2010.

[75] Deborah L Thurston. Real and misconceived limitations to decision based design with utility analysis. *J. Mech. Des.*, 123(2):176–182, 2001.

[76] Matthew G Green, Dan Jensen, Carolyn C Seepersad, and Kristin L Wood. Design for frontier contexts: Classroom assessment of a new design methodology with humanitarian applications. *International Journal of Engineering Education*, 25(5):1029, 2009.

[77] Lin He, Wei Chen, Christopher Hoyle, and Bernard Yannou. Choice modeling for usage context-based design. *J. Mech. Des.*, 134(3), 2012.

[78] Lin He, Wei Chen, and Guenter Conzelmann. On usage context of hybrid electric vehicles in choice studies. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 54822, pages 55–66, 2011.

[79] Horace Herring and Robin Roy. Technological innovation, energy efficient design and the rebound effect. *Technovation*, 27(4):194–203, 2007.

[80] David L Greene. Rebound 2007: Analysis of us light-duty vehicle travel statistics. *Energy Policy*, 41:14–28, 2012.

[81] Morteza Taiebat, Samuel Stolper, and Ming Xu. Forecasting the impact of connected and automated vehicles on energy use: a microeconomic study of induced travel and energy rebound. *Applied Energy*, 247:297–308, 2019.

[82] Tyron Louw, Rafael Goncalves, Guilhermina Torrao, Vishnu Radhakrishnan, Wei Lyu, Pablo Puente Guillen, and Natasha Merat. Do drivers change their manual car-following behaviour after automated car-following? *Cognition, Technology & Work*, pages 1–15, 2020.

[83] Moshe Ben-Akiva and Steven R Lerman. *Discrete choice analysis: theory and application to travel demand*. Transportation Studies, 2018.

[84] Frank S Koppelman and Chandra Bhat. A self instructing course in mode choice modeling: multinomial and nested logit models. 2006.

[85] Joaquim RRA Martins and Andrew B Lambe. Multidisciplinary design optimization: a survey of architectures. *AIAA journal*, 51(9):2049–2075, 2013.

[86] Michael Balmer, Kai Nagel, and Bryan Raney. Large-scale multi-agent simulations for transportation applications. In *Intelligent Transportation Systems*, volume 8, pages 205–221. Taylor & Francis, 2004.

[87] Mojtaba Arezoomand and Jesse Austin-Breneman. Investigating optimal communication frequency in multi-disciplinary engineering teams using multi-agent simulation. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 59186, page V02AT03A018. American Society of Mechanical Engineers, 2019.

[88] Leonard Evans and Michael C Frick. Car mass and fatality risk: has the relationship changed? *American journal of public health*, 84(1):33–36, 1994.

[89] Reza Tolouei and Helena Titheridge. Vehicle mass as a determinant of fuel consumption and secondary safety performance. *Transportation research part D: transport and environment*, 14(6):385–399, 2009.

[90] Adrian Lund. The relative safety of large and small passenger vehicles. In *NHTSA Mass-Size Safety Symposium Washington DC, February*, volume 25, 2011.

[91] Jeremy Dick, Elizabeth Hull, and Ken Jackson. *Requirements engineering*. Springer, 2017.

[92] George A Hazelrigg. *Fundamentals of decision making for engineering design and systems engineering*. 2012.

[93] Yung C Shin and Chengying Xu. *Intelligent systems: modeling, optimization, and control*. CRC press, 2017.

[94] David Sousa Nunes, Pei Zhang, and Jorge Sá Silva. A survey on human-in-the-loop applications towards an internet of all. *IEEE Communications Surveys & Tutorials*, 17(2):944–965, 2015.

[95] G.E. Dieter and L.C. Schmidt. *Engineering Design*. McGraw-Hill, 2013.

[96] Mariëlle Geusen, Erik Jan Hultink, and Katrin Eling. Choice of consumer research methods in the front end of new product development. *International Journal of Market Research*, 55(1):81–104, 2013.

[97] Deborah L Thurston. Multiattribute utility analysis in design management. *IEEE Transactions on Engineering Management*, 37(4):296–301, 1990.

[98] Deborah L Thurston. A formal method for subjective design evaluation with multiple attributes. *Research in engineering Design*, 3(2):105–122, 1991.

[99] Hui Li and Shapour Azarm. Product design selection under uncertainty and with competitive advantage. *J. Mech. Des.*, 122(4):411–418, 1999.

[100] Hui Li and Shapour Azarm. An approach for product line design selection under uncertainty and competition. *J. Mech. Des.*, 124(3):385–392, 2002.

[101] Jungmok Ma and Harrison M Kim. Continuous preference trend mining for optimal product design with multiple profit cycles. *Journal of Mechanical Design*, 136(6):061002, 2014.

[102] Conrad S Tucker and Harrison M Kim. Optimal product portfolio formulation by merging predictive data mining with multilevel optimization. *Journal of Mechanical Design*, 130(4):041103, 2008.

[103] Jungmok Ma, Minjung Kwak, and Harrison M Kim. Demand trend mining for predictive life cycle design. *Journal of Cleaner Production*, 68:189–199, 2014.

[104] Morgan WP Tackett, Christopher A Mattson, and Scott M Ferguson. A model for quantifying system evolvability based on excess and capacity. *Journal of Mechanical Design*, 136(5), 2014.

[105] James S Dyer. Maut—multiattribute utility theory. In *Multiple criteria decision analysis: state of the art surveys*, pages 265–292. Springer, 2005.

[106] Moshe E Ben-Akiva, Steven R Lerman, and Steven R Lerman. *Discrete choice analysis: theory and application to travel demand*, volume 9. MIT press, 1985.

[107] Erin Faith MacDonald, Richard Gonzalez, and Panos Papalambros. The construction of preferences for crux and sentinel product attributes. *Journal of Engineering Design*, 20(6):609–626, 2009.

[108] Maude Josée Blondin, Javier Sanchis Sáez, and Panos M Pardalos. Control engineering from classical to intelligent control theory—an overview. In *Computational Intelligence and Optimization Methods for Control Engineering*, pages 1–30. Springer, 2019.

[109] Richard S Sutton, Andrew G Barto, and Ronald J Williams. Reinforcement learning is direct adaptive optimal control. *IEEE Control Systems Magazine*, 12(2):19–22, 1992.

[110] Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. *An introduction to statistical learning*, volume 112. Springer, 2013.

[111] Theodoros Evgeniou, Constantinos Boussios, and Giorgos Zacharia. Generalized robust conjoint estimation. *Marketing Science*, 24(3):415–429, 2005.

[112] Noah Tovares, Peter Boatwright, and Jonathan Cagan. Experiential conjoint analysis: an experience-based method for eliciting, capturing, and modeling consumer preference. *Journal of Mechanical Design*, 136(10):101404, 2014.

[113] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[114] Burr Settles. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2009.

[115] Max K Shepherd and Elliott J Rouse. The vspa foot: a quasi-passive ankle-foot prosthesis with continuously variable stiffness. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 25(12):2375–2386, 2017.

[116] Max K Shepherd, Alejandro F Azocar, Matthew J Major, and Elliott J Rouse. The difference threshold of ankle-foot prosthesis stiffness for persons with transtibial amputation. In *2018 7th IEEE International Conference on Biomedical Robotics and Biomechatronics (Biorob)*, pages 100–104. IEEE, 2018.

[117] Maysam F. Abbod, Derek A. Linkens, Mahdi Mahfouf, and George Dounias. Survey on the use of smart and adaptive engineering systems in medicine. *Artificial Intelligence in Medicine*, 26(3):179–209, 2002.

[118] Haidawati Nasir, Wan Basyar Wan Aziz, Fuead Ali, Kushsairy Kadir, and Sheroz Khan. The implementation of iot based smart refrigerator system. In *2018 2nd International Conference on Smart Sensors and Application (ICSSA)*, pages 48–52. IEEE, 2018.

[119] Jim Torresen. Reconfigurable logic applied for designing adaptive hardware systems. In *Proc. of the International Conference on Advances in Infrastructure for e-Business, e-Education, e-Science, and e-Medicine on the Internet (SSGRR'2002W)*. Citeseer, 2002.

[120] Alwyn Burger, Christopher Cichiwskyj, Stephan Schmeißer, and Gregor Schiele. The elastic internet of things-a platform for self-integrating and self-adaptive iot-systems with support for embedded adaptive hardware. *Future Generation Computer Systems*, 113:607–619, 2020.

[121] Christian Krupitzer, Felix Maximilian Roth, Sebastian VanSyckel, Gregor Schiele, and Christian Becker. A survey on engineering approaches for self-adaptive systems. *Pervasive and Mobile Computing*, 17:184–206, 2015.

[122] Vikas N O'Reilly-Shah. Factors influencing healthcare provider respondent fatigue answering a globally administered in-app survey. *PeerJ*, 5:e3785, 2017.

[123] Darius Braziunas. Computational approaches to preference elicitation. *Department of Computer Science, University of Toronto, Tech. Rep*, page 62, 2006.

[124] Li Chen and Pearl Pu. Survey of preference elicitation methods. Technical report, 2004.

[125] Joel Huber, Dick R Wittink, John A Fiedler, and Richard Miller. The effectiveness of alternative preference elicitation procedures in predicting choice. *Journal of Marketing Research*, 30(1):105–114, 1993.

[126] Mandy Ryan, Diane Alison Scott, C Reeves, A Bate, Edwin Roland Van Teijlingen, Elizabeth M Russell, M Napper, and CM Robb. Eliciting public preferences for healthcare: a systematic review of techniques. *Health technology assessment (Winchester, England)*, 5(5):1–186, 2001.

[127] Roland Helm, Laura Manthey, Armin Scholl, Michael Steiner, et al. Empirical evaluation of preference elicitation techniques from marketing and decision analysis. Technical report, Friedrich-Schiller-Universität Jena, Wirtschaftswissenschaftliche Fakultät, 2003.

[128] Moshe Emanuel Ben-Akiva, Daniel McFadden, Kenneth Train, et al. *Foundations of stated preference elicitation: Consumer behavior and choice-based conjoint analysis*. Now, 2019.

[129] Austin Le, Benjamin H Han, and Joseph J Palamar. When national drug surveys "take too long": An examination of who is at risk for survey fatigue. *Drug and alcohol dependence*, 225:108769, 2021.

[130] Thomas J Whelan. Antecedents of anonymity perceptions in web-based surveys. In *23rd annual meeting of the Society for Industrial and Organizational Psychology, San Francisco, CA*, 2008.

[131] Graham Kalton. Nonsampling errors in panel surveys. *Panel surveys*, pages 249–270, 1989.

[132] Pazit Ben-Nun et al. Respondent fatigue. *Encyclopedia of survey research methods*, 2:742–743, 2008.

[133] Ian Brace. *Questionnaire design: How to plan, structure and write survey material for effective market research*. Kogan Page Publishers, 2018.

[134] Don A Dillman, Glenn Phelps, Robert Tortora, Karen Swift, Julie Kohrell, Jodi Berck, and Benjamin L Messer. Response rate and measurement differences in mixed-mode surveys using mail, telephone, interactive voice response (ivr) and the internet. *Social science research*, 38(1):1–18, 2009.

[135] Stephen R Porter, Michael E Whitcomb, and William H Weitzer. Multiple surveys of students and survey fatigue. *New directions for institutional research*, 2004(121):63–73, 2004.

[136] Barbara Schmidt-Belz. User trust in adaptive systems. In *LWA*, pages 69–73, 2005.

[137] Vanessa Evers, Henriette Cramer, Maarten van Someren, and Bob Wielinga. Interacting with adaptive systems. In *Interactive collaborative information systems*, pages 299–325. Springer, 2010.

[138] Miriam Gil, Vicente Pelechano, Joan Fons, and Manoli Albert. Designing the human in the loop of self-adaptive systems. In *International Conference on Ubiquitous Computing and Ambient Intelligence*, pages 437–449. Springer, 2016.

[139] Jungchul Park and Sung H Han. A fuzzy rule-based approach to modeling affective user satisfaction towards office chair design. *International Journal of Industrial Ergonomics*, 34(1):31–47, 2004.

[140] Myung H Yun, Sung H Han, Taebeum Ryu, and Keumsun Yoo. Determination of critical design variables based on the characteristics of product image/impression: Case study of office chair design. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 45, pages 712–716. SAGE Publications Sage CA: Los Angeles, CA, 2001.

[141] Liesbeth Groenesteijn, Peter Vink, Michiel de Looze, and Frank Krause. Effects of differences in office chair controls, seat and backrest angle design in relation to tasks. *Applied ergonomics*, 40(3):362–370, 2009.

[142] Rensis Likert. A technique for the measurement of attitudes. *Archives of psychology*, 1932.

[143] Henry B Mann and Donald R Whitney. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics*, pages 50–60, 1947.

[144] Constantin Mircioiu and Jeffrey Atkinson. A comparison of parametric and non-parametric methods applied to a likert scale. *Pharmacy*, 5(2):26, 2017.

[145] JFC De Winter and Dimitra Dodou. Five-point likert items: t test versus mann-whitney-wilcoxon (addendum added october 2012). *Practical Assessment, Research, and Evaluation*, 15(1):11, 2010.

[146] Pedro Cosme Vieira. T-test with likert scale variables. *Available at SSRN 2770035*, 2016.

[147] Rudolf N Cardinal and Michael RF Aitken. *ANOVA for the behavioral sciences researcher*. Psychology Press, 2013.

[148] Ron C Mittelhammer, George G Judge, and Douglas J Miller. *Econometric foundations pack with CD-ROM*. Cambridge University Press, 2000.

[149] G Rupert Jr et al. *Simultaneous statistical inference*. Springer Science & Business Media, 2012.

[150] Paul Resnick and Hal R Varian. Recommender systems. *Communications of the ACM*, 40(3):56–58, 1997.

[151] Geoff Boeing. Osmnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems*, 65:126–139, 2017.

[152] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. MIT press, 2009.