ENGINEERING VISUAL DISPLAYS TO INFLUENCE CHOICE IN AUTOMATED DECISION SUPPORT SYSTEMS

A Dissertation
Presented to
The Academic Faculty

by

Cale Michael Darling

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Psychology

Georgia Institute of Technology
August 2018

ENGINEERING VISUAL DISPLAYS TO INFLUENCE CHOICE IN AUTOMATED DECISION SUPPORT SYSTEMS

Approved by:

Dr. Francis T. Durso, Advisor
School of Psychology
Georgia Institute of Technology

Dr. Rustin D. Meyer
Department of Psychology
Penn State University

Dr. Rickey P. Thomas]
School of Psychology
Georgia Institute of Technology

Dr. Jeffrey R. Parker
College of Business
Georgia State University

Dr. Bruce N. Walker
School of Psychology
Georgia Institute of Technology

Date Approved: July 6th, 2018

ACKNOWLEDGEMENTS

I would like to thank my advisor, Frank Durso, for sharing his wisdom and offering his guidance throughout my graduate career. I would also like to thank my committee members, Rustin Meyer, Jeff Parker, Rick Thomas, and Bruce Walker for their insightful feedback which helped me shape this dissertation. I would like to thank Josh Hoffman and Julian Sanchez for encouraging me to push forward with this dissertation while I started my career at Deere & Company. Finally, I would like to especially thank my parents, Vicki and Keith Darling for their love and support, without which this dissertation would not have been possible.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

The task of choosing between decision alternatives presented on a visual display is ubiquitous. Automated decision support systems (DSS) provide a powerful means of improving human decision-making outcomes, but they can also introduce deleterious effects in the form of automation bias (e.g., commission errors and errors of omission). Research has shown that informationally equivalent display designs can lead to significant differences in terms of decision-making outcomes. The current study examined how the influence of visual display design factors on decision making can be leveraged to increase compliance with an automated DSS and reduce potential automation bias. To this end, a series of four experiments were conducted. In each experiment, participants completed a simulated route navigation task in which they were tasked with choosing one of four different routes that were described by four different attributes in order to navigate to their destination. Experiments 1 and 2 were designed to explore how display design factors could be used to establish a decision environment that influenced participants choices in a predictable manner. Results revealed that highlighting an attribute in yellow to increase its perceptual salience increased the likelihood that participants would choose the route that was strongest on the salient attribute. Experiments 3 and 4 applied this salience effect to the design of an automated DSS which recommended one of the four routes to participants. By increasing the salience of an attribute, choice share in favor of the route recommended by the automated DSS increased by as high as 15%. However, this increase in compliance came at the cost of increasing commission errors; participants chose the recommended route even on trials in which it was inferior to other alternatives. For this reason, salience

effects should be applied to cases in which the cost of commission errors is low or when automation reliability is high. Informationally equivalent display design factors can be manipulated to increase compliance, but to reduce automation bias, the display design must communicate the logic underlying the automation's recommendations.

CHAPTER 1. INTRODUCTION

The task of choosing between decision alternatives presented on a visual display is ubiquitous. Consider for example, a global positioning system (GPS) providing a series of potential routes for you to take to a new restaurant. Upon arriving at the restaurant, you view a food menu displaying several dishes, from which you choose one. Choice decision-making tasks such as these are fairly common examples. However, choice tasks can range from the routine to more extreme, safety-critical environments such as a physician choosing a treatment plan from a set of pre-selected options to administer to a patient. In these examples, changing the manner in which decision information is displayed presents an opportunity to help people make better decisions.

The way in which decision information is displayed can significantly influence how people choose among alternatives. Visual displays represent the medium through which decision information is communicated to, and subsequently processed by the decision maker. Research has shown that the same information, when presented using different display designs, can lead to significant differences in terms of decision-making outcomes (see Hegarty, 2011). Explaining the effect of display design on decision making, cognitive scientists have argued that informationally equivalent displays (those which contain the same information) are not necessarily computationally equivalent (Larkin & Simon, 1987). In other words, visual displays can be manipulated independent of the decision information that is being presented. For this reason, visual display design constitutes an important factor for understanding task performance in general, and decision-making outcomes in

particular. Understanding display design effects on decision making is an important part of designing a beneficial automated decision support system.

In an effort to help people make better decisions, automated decision support systems can aid the human decision-making process. In general, automated decision support systems (DSS) present a set of decision alternatives, from which the human can choose. By leveraging the computational capabilities of automation, the DSS can examine more alternatives in less time than their human counterpart and recommend decision alternatives without restricting choice. However, imperfect automation can also introduce deleterious effects in the form of decision-making biases that arise from over or under relying on the automation (Goddard, Roudsari, & Wyatt, 2011). Therefore, leveraging display design effects on decision making offers the potential to influence how people use such automated DSS to ensure that people get the most out of such systems.

The visual display is often the primary means through which such highly automated systems communicate decision information to the human. Highly automated systems are growing in ubiquity and are now being used in more safety-critical environments by well-trained operators as well as by novice consumers at home and on the road (West, 2015). Designing for interactions with highly automated systems will become a more common, yet critical human factors challenge faced by designers across industries (Hancock, 2017). To this point, researchers have recently called for a paradigm shift in how we design for interacting with highly automated systems (de Visser, Pak, & Shaw, 2018). In the progression from automation to autonomy, the task of choosing among different plans of action for a highly automated system will likely become a more dominant paradigm of

interacting with technology. Therefore, understanding how to engineer a decision environment to influence choice can be directly applied to improve peoples' interactions with highly automated systems.

Seemingly insignificant visual display design factors can be leveraged to strategically influence peoples' decision-making processes and ultimately their choice of alternatives. For this reason, visual display and graphical user interface (GUI) designers are uniquely positioned to help people make better decisions. There is almost no end to how a visual display can be manipulated. One can easily imagine how decision information could be displayed in a manner that all but omits decision-makers from considering one or more alternatives (e.g., the use of a micro font that significantly reduces legibility). On the other hand, subtle display factors such as the order in which alternatives are sequenced within the display can impact how people evaluate and choose an alternative. When considering how to engineer displays to influence choice, it is important that the user not be overly restricted in their ability to choose a given option, for doing so could instill distrust and contribute to disuse of the system (Szalma, 2014). Thus, rather than identify display design factors that influence choice at all costs, the current dissertation focused on identifying those display design factors that retain information equivalence but do not restrict the decision-maker's ability to choose. Such display elements are an important consideration for the design of an automated DSS that encourages appropriate use of the automated decision aid.

The main goal of this dissertation was to understand how to design a decision environment in order to influence how people used an automated DSS. To this end,

research at the intersection of visual display design, choice decision-making, and human-automation interaction was reviewed. From this review, four fundamental characteristics of information displays were identified, each of which can be manipulated to influence decision making without displaying additional information to the user (i.e., retaining informational equivalence). By leveraging the predictable effects of display factors in the design of an automated DSS, the current study examined how automation designers can influence people to comply with the DSS when they otherwise might not have. Furthermore, going beyond influencing choice and increasing compliance, the current study explored how display design factors might be used to ultimately foster appropriate use of an automated DSS.

1.1 Human-Automation Interaction and Automation Bias

Automated DSS constitute a powerful method of improving human decision-making. The enhanced computational abilities of automated systems allow for simultaneously evaluating numerous decision alternatives in a matter of milliseconds and recommending a subset of those alternatives to the human. Depending upon its level of reliability, automation can significantly improve decision-making performance over humans alone (Goddard et al., 2011). For instance, GPS navigation systems contain an automated decision aid that evaluates multiple routes to a given destination and offers recommendations to the driver. Moreover, DSS can reduce the prevalence of decision-making biases that arise from the use of heuristic-based decision making and which result in poor choices or sub-optimal decisions (Kahneman & Frederick, 2002; Kahneman &

Tversky, 1979). Thus, automated DSS provide a way to influence peoples choices and help them make better decisions.

Although automated DSS can improve decision-making performance, imperfect automation can be notorious for introducing new performance concerns, one of which is the phenomenon of automation bias (Goddard et al., 2014; Mosier, Skitka, Heers, & Burdick, 1998). In general, automation bias results from operators using the outcome of a decision aid as a heuristic, which replaces what would otherwise involve an effortful process of extracting and evaluating information required to reach a decision (Mosier & Skitka, 1996). The key distinction between automation-induced bias and well-established decision-making biases found in the decision-making literature is that the biases arise from the operator's interaction with the automated system (Parasuraman & Manzey, 2010). For example, as an operator interacts with an automated DSS, it can lead to commission errors, which occur when people rely on the automation's recommendations despite it being inaccurate or even though there is a better alternative readily available.

The current study focused on designing a decision environment for decision support as opposed to supervisory control. However, researchers distinguish between supervisory control and decision support as two distinct categories of HAI paradigms in order to systematically study automation-induced biases (Parasuraman & Manzey, 2010). The supervisory control paradigm typically involves an operator using an automated monitoring system to perform routine checks on system states and alert the operator when a critical system state is immanent. An example of such supervisory control interactions is the modern air traffic controller using an automated aircraft conflict detection system. In

contrast, the decision support paradigm involves an automated decision aid that supplements the operator's cognitive processes. Types of decision support interactions can range from simply cueing the operator to attend to specific information, to recommending and implementing a course of action (Parasuraman, Sheridan, & Wickens, 2000). For example, in-vehicle navigation systems provide a route and inform the driver when to turn at an intersection along that route.

With each interaction paradigm there are different types of automation-induced biases that can lead to performance decrements. For the supervisory control paradigm, the main performance issue arises due to automation-induced complacency, which is defined as "poorer detection of system malfunctions under automation compared with under manual control" (Parasuraman & Manzey, 2010, p. 390). In contrast, the decision support paradigm leads to automation bias, the most relevant of which is a commission error. In the HAI literature, some researchers use the term automation bias to refer to any automation-induced bias. For the sake of clarity and given the focus on DSS in the current paper, automation bias is used to refer specifically to commission errors and errors of omission.

Within a DSS, commission errors occur when people rely on the automation's recommendations rather than examine the decision information and critically evaluate the automation's recommendations (Parasuraman & Manzey, 2010). Researchers have proposed that one of the main factors contributing to the occurrence of commission errors is the human tendency to conserve cognitive effort (Mosier & Skitka, 1996; Wickens & Hollands, 2000). In a desire to conserve effort, people avoid the effortful process of

evaluating the available decision information and instead they willfully accept the automation's recommendation regardless of its validity. Of course, compliance with the automation can lead to a positive outcome, so long as the automation is perfect; however, such perfectly reliable systems are nearly non-existent.

In designing an automated DSS, manipulating the way that decision information and recommended alternatives are displayed presents an opportunity for improving HAI outcomes and reducing automation bias. Display design can significantly impact how people perceive and interact with automated systems; after all, "because direct observation of the automation is often impractical or impossible, perception of the automation-related information is usually mediated by a display" (Lee & See, 2004, p. 25). Leveraging display factors that influence choice in a predictable fashion, designers can directly address issues of increasing compliance and reducing automation bias. Recognizing the importance of display design for improving HAI outcomes, researchers have explored the notion of increasing automation transparency.

1.2 Designing for Transparency in Automated Decision Support Systems

Automation transparency is the extent to which the logic underlying an automated system is made available to the human operator (Seong & Bisantz, 2008). The level of transparency can be increased through display design in numerous ways including feedback for assessing performance and cueing people to attend to specific information. This latter approach to increasing automation transparency represents a direct application of using display design factors to influence decision-making processes. For example,

highlighting information can be used to capture an operator's selective attention and help convey why an automated system is recommending one decision alternative over another. If an operator understands the logic underlying the automation's recommendation, then they are more likely to accept the recommendation when it is valid and ignore it when it is not. Thus, it is tenable that visual display factors like highlighting can be used to increase transparency; in turn, reducing automation bias and encouraging appropriate use of automation.

Researchers have decomposed the concept of automation transparency, proposed methods of designing for transparency, and examined its effects on HAI outcomes (Chen et al., 2014; Lyons, 2013; Ososky, Sanders, Jentsch, Hancock, & Chen, 2014). For example, Chen and colleagues (2014) proposed a situation awareness-based agent transparency (SAT) model that outlines transparency requirements at three different levels, which correspond with Endsley's (1995) model of SA. At level 1, transparency should help the operator understand what the system is doing by providing basic information about the system's current state, goals and planned actions. For level 2, operators should understand why the system is doing what it is doing by displaying information that establishes the rationale or logic underlying the system's actions. At level 3, operators are provided with information regarding the projected future state of the system along with information regarding the anticipated consequences, probability of success, and level of certainty. Not all three levels are necessary in order to increase automation transparency; instead, each level simply represents a different aspect of the concept which can be designed for

accordingly. With this structured approach to transparency, the SAT model can be applied to user interface design to increase transparency and reduce automation bias.

Increasing automation transparency through display design of an automated DSS can help reduce automation bias (Mercado et al., 2016; Rovira, Cross, Leitch, & Bonaceto, 2014; Sadler et al., 2016; Stowers et al., 2016; Wright, Chen, Barnes, & Hancock, 2016a). If a person understands why the automated DSS has recommended a particular alternative (via increased automation transparency) then they are less likely to fall victim to automation bias and commission errors in particular. These studies involve similar versions of an automated decision support paradigm in which participants were tasked with managing unoccupied vehicles in varying work domains including air traffic control and military surveillance. Furthermore, transparency was manipulated by providing operators with additional information to convey the rationale regarding the system's recommendations or provide a statement of risk associated with each recommendation.

For example, Mercado et al. (2016) had participants manage multiple unoccupied vehicles and they were tasked with evaluating different plans of action for the vehicles. An automated decision aid supported participants' decision-making by recommending one of two different plans. Participants were responsible for choosing a plan for the system to implement. There were three levels of automation transparency which correspond with the three levels in Chen et al.'s SAT model: Level 1 displayed basic information about the vehicles to be used in that plan; level 1+2 included the basic information plus information to establish the rationale underlying the system's recommended plan (e.g., "arriving to destination faster with adequate coverage"); level 1+2+3 included basic information, the

rationale statement, and introduced a colored icon to convey the system's level of uncertainty for that plan's information. Automation reliability was fixed at 70%; thus, the recommended plan was inferior on 30% of trials. Analysis of plan selection data showed that the percentage of correct choices, both in terms of correct acceptance and correct rejection of the recommended plans increased as a function of the level of transparency. Thus, increasing transparency significantly reduced automation bias that might have otherwise led to misuse of the automation.

To explain how increasing transparency can reduce automation bias, researchers have examined the relationship between transparency and trust in automation (Lyons et al., 2017; Lyons et al., 2016; Rovira et al., 2014; Mercado et al., 2016; Sadler et al., 2016; Selkowitz, Lakhmani, Chen, & Boyce, 2015). For instance, Lyons and colleagues (2016) manipulated the level of transparency associated with an automated DSS's recommended airport runway diversion and found that increasing levels of transparency led to a significant increase in participants' trust in the automation. However, simply increasing trust in the automation does not translate to improved performance if the automation is imperfect.

With imperfect automation, it is more important for designers to find ways to foster appropriate trust rather than pursue a general increase in trust levels (Lee & See, 2004). Blindly trusting the automation will lead to complacency and overreliance on the automation; therefore, display design factors intended to increase transparency need to help the human understand when they should trust the system, and thus foster appropriate trust in the automation. Indeed, some evidence suggests that increasing transparency can

increase trust, but further increasing the level of transparency does not necessarily lead to further increases in trust (Mercado et al., 2016; Wright et al., 2016a). Such results indicate that increasing transparency can help people understand why the system is recommending an alternative, in turn enabling people to better gauge when they should trust the automation.

The majority of research demonstrating how transparency can reduce automation bias involves displaying additional information to the human. For example, to manipulate transparency of an automated DSS for managing unoccupied vehicles, Mercado et al. (2016) included a statement summarizing how the recommended plan compared to others on task-specific variables (e.g., "arriving to destination faster with adequate coverage"), in addition to basic information about the plan. Lyons et al. (2016; 2017) manipulated transparency by including a statement of risk associated with recommended airport runway diversions. Of course, these are valid approaches to increasing transparency; however, it is important to consider that increasing the amount of information might impact operator workload and increase the time taken to process and evaluate the decision alternatives.

Research has shown that a moderate level of workload is optimal for helping operators maintain situation awareness and sustaining task performance levels (Durso & Alexander, 2010; Lee & See, 2004; Parasuraman & Wickens, 2008). Therefore, it is important to consider how different methods of increasing transparency might increase operator workload. For example, including an additional statement summarizing the risk associated with each plan recommended by the automation might increase workload, reduce situation awareness, and lead to performance decrements for concurrent tasks.

Several studies have investigated the effects of increasing transparency on operator workload, but the results are mixed. Some researchers have found that displaying additional information to increase transparency can increase subjective workload (Dorneich et al., 2017; Rovira et al., 2014) as well as objective measures of workload using eye tracking data (Wright et al., 2016b). Furthermore, there is evidence to suggest that increasing transparency can lead to increased decision response times (Helldin, 2014; Rovira et al., 2014; Wright et al., 2016b) which can indicate increased information processing demands and workload. However, other studies have used similar methods of increasing transparency yet found no effect on subjective workload (Selkowitz et al., 2015; Stowers et al., 2017), nor objective workload (Mercado et al., 2016). Suffice to say, the relationship between transparency and workload is complex. However, it is possible that displaying additional information to increase transparency might reduce automation bias at the cost of increasing workload.

In comparison to the reviewed methods of designing for transparency, visual display designs that retain informational equivalence might provide a subtler means of influencing choice, increasing transparency, and reducing automation bias. The reviewed research on automation transparency has shown how displaying information that conveys the rationale and logic underlying an automated DSS can influence choice and reduce automation bias. Some researchers have argued that increasing transparency requires displaying additional information (Chen et al., 2014). Challenging this notion, the current study explored how strategically manipulating fundamental characteristics of a decision environment might influence choice and increase transparency of an automated DSS.

1.3 Visual Display Design and Decision Making

Previous research has shown that the same information presented using a different display design can lead to significant differences in terms of decision-making outcomes (Hegarty, 2011). By leveraging the influence of visual display design on choice, automation designers can increase transparency without displaying additional information, which might otherwise increase decision-makers workload and time to choose. In the following subsections, an overview of the choice decision-making process is presented to establish how display design can systematically influence choice. Next, research investigating the effects of display design factors on decision making from several disciplines is briefly reviewed to establish how the current dissertation differs from other approaches. Finally, this section concludes with a summary and an introduction to the decision environment that was used in the current study.

1.3.1 Overview of the Choice Decision-Making Process

A conceptual understanding of the choice decision-making process is necessary to appreciate how display design can influence choice. Researchers have proposed that in general, people often approach a decision without clearly defined preferences for one alternative (Bettman, Luce, & Payne, 1998). Instead, people construct their preferences as they evaluate alternatives. This preference construction process can be depicted by multiple, interactive stages that represent decision processes.

For Payne, Bettman, and Schkade (2000), the decision-making process involves four main stages which are depicted in Figure 1. First, people form an initial cognitive

representation of the decision problem as they recognize the decision space and begin forming expectations of the alternatives. Next, the decision maker allocates their selective attention in order to acquire and interpret the displayed decision information (e.g., identifying different attribute values and retaining them in working memory). The interpreted information is internally processed by combining different attributes and associating them with their respective alternatives. The information is then evaluated until one alternative exceeds a choice threshold and finally the decision-maker articulates their choice.

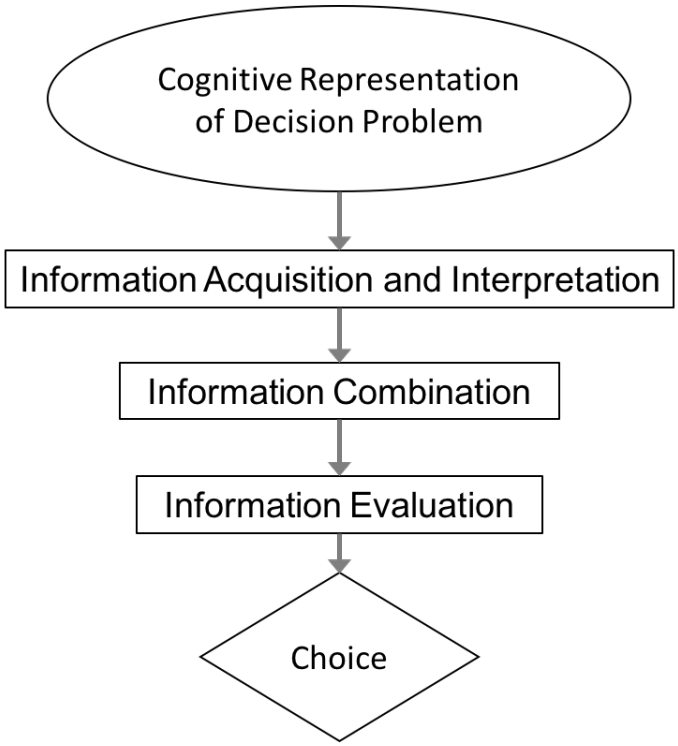

Figure 1 – Conceptual model of the choice decision-making process recreated from Payne et al. (2000). Depicts multiple interacting stages of decision-making to capture how preferences are constructed.

Decision makers adopt various decision strategies (e.g., heuristics) to evaluate alternatives, construct preferences, and choose an alternative (Bettman et al., 1998; Bettman, Johnson, & Payne, 1991; Payne, Bettman, & Johnson, 1993; Payne et al., 2000). Bettman and colleagues (1998) proposed that information processing limitations associated with bounded rationality (Simon, 1955) lead people to adopt decision strategies that serve to facilitate the acquisition and evaluation of decision information. In effect, adopting decision strategies can reduce the cognitive effort associated with a particular decision process (e.g., information combination).

Heuristic-based decision making can help people reduce cognitive effort in different ways, but it can also systematically bias how people evaluate decision information (Shah & Oppenheimer, 2008). Shah and Oppenheimer identified five methods for effort-reduction: 1) examining fewer cues, 2) reducing the difficulty associated with retrieving and storing cue values, 3) simplifying the weighting principle for cues, 4) integrating less information, 5) examining fewer alternatives. According to Shah and Oppenheimer, every decision-making heuristic involves one or more of these effort-reduction methods. The effort-reduction methods that people use in a given situation can shape how people evaluate alternatives and systematically bias choice in a predictable fashion. Furthermore, the methods people use in a particular situation will depend upon several factors, including the way the information is displayed.

The manner in which decision information is displayed can determine the amount of cognitive effort that is required to evaluate and choose an alternative (Hegarty, 2011; Shah & Oppenheimer, 2009). If two informationally equivalent displays are not

computationally equivalent, then the display design will determine what method of effort-reduction people use. If the display design leads people to use a consistent method of effort-reduction, then their choice among alternatives will be systematically biased in a predictable manner. Furthermore, the concreteness principle (Slovic, 1972) proposes that people tend to use only the displayed information and they use it in the way in which it is displayed, as opposed to converting it to better fit the task. Therefore, people are unlikely to restructure the information and instead conform to a given method of effort-reduction. For example, consider a list of decision alternatives described by multiple attributes. If the alternatives are sequenced by descending value on a "product quality" attribute, then the display design will lead people to simplify how attributes are weighted, thus increasing the likelihood that people will choose an alternative that is high on the product quality attribute.

To summarize, display design can determine how people will conserve cognitive effort, which can systematically bias peoples' decision-making process in favor of a particular alternative. Considering the context of an automated DSS, display design effects on choice can be leveraged to increase compliance and potentially reduce automation bias.

1.3.2 Designing a Decision Environment to Influence Choice

The influence of visual display design on choice has been researched across several disciplines including human factors, judgment and decision making, marketing, and consumer decision making. Within the field of human factors, the primary focus has been on improving task performance by increasing the legibility and usability of displayed information. Research from the consumer decision-making literature provides insights into

specific display design factors that have been shown to influence how consumers choose among products or decision alternatives. Furthermore, recognizing that the manner in which decision information is presented to people can significantly influence their decision-making behavior has led to the emergence of the field of choice architecture.

Choice architecture studies how the manner in which choices are presented to people can significantly change their decision making and behavior (Johnson et al., 2012). The term choice architecture was coined by Thaler and Sunstein in reference to how the presentation of decision information provides the opportunity to nudge or "alter people's behavior in a predictable way" (Thaler & Sunstein, 2008, p. 6). A nudge refers to an aspect of the decision environment that influences peoples' choices and behavior (Thaler, Sunstein, & Balz, 2012). Choice architecture emerged from the application of behavioral economics with the goal of nudging people to make personally and socially beneficial choices (Thaler et al., 2012). A key assumption underlying choice architecture is that humans often act irrationally due to cognitive limitations associated with bounded rationality (Simon, 1955). Therefore, decision making is often biased in predictable ways that lead to suboptimal choices. Choice architecture proposes nudges or ways to influence choice by leveraging known decision-making biases in the design of the decision environment.

Choice architecture offers methods of designing a decision environment to nudge people towards beneficial choices. The tools or methods of choice architecture can be grouped into those that involve setting up the task and those that describe the options (Johnson et al., 2012; Münscher, Vetter, & Scheuerle, 2015). Methods used for setting up

the task include reducing the number of alternatives and using default options. These methods primarily involve changing the decision information or altering the procedure itself. In comparison, methods for describing the options include partitioning options and rescaling attribute values to highlight differences between alternatives. Both categories constitute important aspects of designing a decision environment; however, many of the choice architect's methods involve changing the decision information itself. The result is a display that is not informationally equivalent, or it impacts the decision-maker's ability to choose, like in the case of the default options method.

The main goal of the current dissertation was to create a decision environment that increased compliance with an automated DSS. Clearly, this goal falls under the umbrella of choice architecture in that designing a DSS involves designing a choice environment. However, the aim of choice architecture is broader in scope; choice architecture goes beyond identifying informationally equivalent display design factors that subtly influence choice. By focusing on visual display factors that retain informational equivalence, the current dissertation investigated aspects of the display design that have received less attention in the choice architecture literature. Regardless, the results of the current study might offer another set of tools for the choice architect.

In terms of systematically examining how subtle display design factors influence choice, the consumer decision-making literature offers the most support. There are many ways to manipulate characteristics of visual displays to influence choice without introducing additional information. Kleinmuntz and Schkade (1993) identified three component characteristics of information displays that can impact decision-making

processes: Format, organization, and sequence. These three categories represent "fundamental characteristics of displays that one can, in principle, vary independently" (Kleinmuntz & Schkade, 1993, p. 222). In addition, I have identified a fourth category, referred to as information salience, which constitutes any manipulation to the perceptual salience of a specific piece of decision information as it is displayed to the decision maker. Research from the consumer decision-making literature has examined how manipulating factors from these categories can influence decision-making processes and ultimately choice among products. Consumer decision-making tasks often involve displaying several alternatives in a multi-attribute decision-making task, which is comparable to the decision-making paradigm used in the current study.

Designing a decision environment for the context of a multi-attribute choice decision task requires instantiating format, organization, sequence, and salience. Extant research regarding the influence of factors from each of the four categories was reviewed. By identifying how each factor impacted decision-making processes and influenced choice, previous research was used to inform the design of the decision environment used in the current study (see Appendix A for a literature review of each factor).

1.4 The Current Study

The current study investigated how to engineer the visual display of decision information to increase compliance with an automated DSS. To this end, a series of four experiments was conducted. Experiments 1 and 2 were designed to address open questions regarding how sequence and salience display factors should be manipulated. With an

understanding of how to design a decision environment that influences choice, experiments 3 and 4 demonstrated how these display design effects can be used to increase compliance with an automated DSS.

In the current study, the decision-making task was situated in the context of a simulated vehicle navigation task. Participants were responsible for choosing from a set of four different routes in order to navigate to their destination. Routes constituted decision alternatives and each route was described by four attributes. Visual display factors were manipulated to examine how they might influence choice in favor of a targeted route.

The route displays were presented in a table format and organized by alternative. Each route (i.e., decision alternative) was presented as a table, which contained the route name and its four attributes. Tables have been shown to influence choice by encouraging people to selectively process decision information as opposed to comprehensively evaluating each decision alternative (Dilla & Steinbart, 2005). Furthermore, by facilitating selective processing, tables allow other display factors to exert influence on choice, such as information sequence and salience. Regarding organization, the route displays were organized by alternative as opposed to by attribute. That is, each route was displayed as a separate table on a separate screen. Using a table format that is organized by alternatives can lead people to selectively process the decision alternatives, which can bolster potential sequence and salience effects.

To complete the design of the decision environment, information sequence and salience were instantiated. Although there is evidence that manipulating information

sequence and salience factors can influence choice (e.g., Jiang & Punj, 2010), it is unclear how both factors can be combined to strengthen their influence or if one factor is sufficient for influencing choice. To address this limitation, it was necessary to conduct two initial experiments.

The sequence of alternatives and the salience of attributes were manipulated in experiments 1 and 2 to determine how each factor should be instantiated in the design of the automated DSS. The sequence manipulation determined the order in which alternatives were presented to participants. Specifically, route displays were presented in either descending attribute value or in a randomized order. Previous research suggests that ordering alternatives by descending attribute values can influence how people weigh the attributes, and thus increase choice in favor of an alternative that is higher on the sequenced attribute (e.g., Cai & Xu, 2008). The salience factor was manipulated at the attribute level by highlighting one of four attributes in yellow or with no highlighting to serve as a control. The rationale for highlighting an attribute was to capture participants' selective attention, increase the weight given to that attribute, and bias evaluations of alternatives based on their highlighted attribute values. Manipulating the sequence of alternatives and salience of attributes provided a means of designing the decision environment that could influence choice in favor of a specific, targeted alternative.

In summary, the goal of this dissertation was to design a decision environment that influenced choice of alternatives recommended by an automated DSS. In order to make an informed decision about how each factor should be configured in the automated DSS's visual display, two experiments were conducted to investigate how sequence and salience

should be manipulated. Experiments 1 and 2 investigated the sequence and salience effects within the described decision environment and the results were used to inform the display design in experiments 3 and 4. Information format and organization were controlled across all four experiments, but both were instantiated based on how they facilitated sequence and salience effects. Therefore, the resulting display design allowed for examining sequence and salience effects, but it ultimately served as an important first step toward the goal of designing a decision environment that influenced choice. Armed with an understanding of how to instantiate sequence and salience, experiments 3 and 4 investigated how to leverage display design effects to increase compliance with an automated DSS.

CHAPTER 2. EXPERIMENT 1

In this experiment, the sequence of alternatives and the salience of attributes was manipulated to inform how each factor should be manipulated in the design of an automated DSS. A significant sequence effect was predicted in that when alternatives were sequenced by descending value on a targeted attribute, choice share for the route highest on the targeted attribute would increase. Similarly, salience was predicted to increase choice share for the route highest on the targeted, salient attribute. Finally, a sequence by salience interaction was predicted in that combining sequence and salience effects would bolster the relative influence of each factor. In other words, the choice share of the targeted alternative would be higher than when either sequence or salience was manipulated in isolation.

2.1 Method

2.1.1 Participants

A total of 76 undergraduates from the Georgia Institute of Technology participated in this experiment. Participants were enrolled in a psychology course and elected to participate in the study by selecting it from a list of experiments displayed in the SONA online experiment scheduler. The approximate 30-minute experiment was completed in partial fulfillment of a research familiarization requirement.

2.1.2 Materials

The stimuli were 16 computer-generated tables, each of which constituted one route display (i.e., decision alternative). Figure 2 provides an example of a route display which includes the route name, separate columns for each attribute label, and the corresponding attribute value on the bottom row of the table. All text in the route display was presented in a 12-point Arial font. Each route display was given a unique name so that participants could distinguish each route within the set. Routes were named by randomly selecting street names from the United States using a list randomization tool (www.randomlists.com).

Main Street			
Road Quality	**Speed**	**Fuel Efficiency**	**Traffic Flow**
Very Low	Low	Very High	High

Main Street			
Road Quality	**Speed**	**Fuel Efficiency**	**Traffic Flow**
Very Low	Low	Very High	High

Figure 2 – Example of a route display with no salience manipulation (top) and the same route display with the salience manipulation (bottom). Each route display included the name of the route, the four attribute labels and their corresponding values for that route. In this example, Road Quality constitutes the targeted attribute.

On each trial, participants chose their preferred route from a set of four different routes (referred to as a route set) prior to navigating to their destination and delivering each package. All routes were described by four component attributes, which included road quality, speed, traffic flow, and fuel efficiency. Attribute values varied along a five-point scale with verbal descriptions for each discrete value (1 = “Very Low”, 5 = “Very High”);

higher values indicated positive valence (e.g., very high traffic flow indicated that traffic was moving quickly along that route).

Each route's overall attribute utility value was calculated by taking the sum of that route's four attribute values. The highest utility value that a given route could have was 20 (i.e., "Very High" for all four attributes) and the lowest was 4 (i.e., all four attribute values were "Very Low"). In this experiment, the overall attribute utility value ranged from 11 to 16. A list of each route's attribute values and the overall utility is available in Appendix E. Of the 16 trials, four involved a route utility value of 11, eight trials with utility value of 13, and four trials with utility value of 16.

The overall attribute utility value for each route was controlled such that no route dominated the others within a given route set. For example, Route A might be dominant on fuel efficiency, but its remaining attribute values were instantiated such that its overall utility did not exceed that of the three remaining routes in its route set. Therefore, a given route might be dominant on one attribute, but no route was dominant when considering the overall attribute utility value of each route.

The four route attributes were selected based on relevance to the vehicle navigation task and because each could be informative in terms of one's preferences for a given route. Furthermore, the attributes were chosen in order to exclude attributes which might be overly diagnostic when evaluating a given route set. For instance, estimated trip duration was excluded from the set of attributes because it is arguably the most important attribute that people use when choosing among routes provided by a GPS roadway navigation

system. Excluding such attributes was intended to reduce the likelihood that participants would choose a route based solely on a single attribute, regardless of how display factors were manipulated.

The GPS navigation animation constituted a low-fidelity, generic map that depicted the participant's vehicle location as it traveled along the route to the destination for package delivery. A screenshot of the navigation animation is presented in Figure 3. The animations were designed to be similar to GPS roadway navigation commonly used on smartphone apps and in-vehicle displays. All GPS animations were created using the graphic design software, Sketch 3 (Bohemian Coding, The Hague, Netherlands) and exported to an animated graphic interchange format (GIF) file that was embedded within a survey page.

The purpose of the navigation animation was to simulate a route monitoring task and establish a context that motivated participants choice of routes. By animating the route navigation, the intention was to create an experience in which the participants choice of routes had a direct impact on how they completed the task. The intended result was to provide motivation for participants to choose a route which they might otherwise lack.

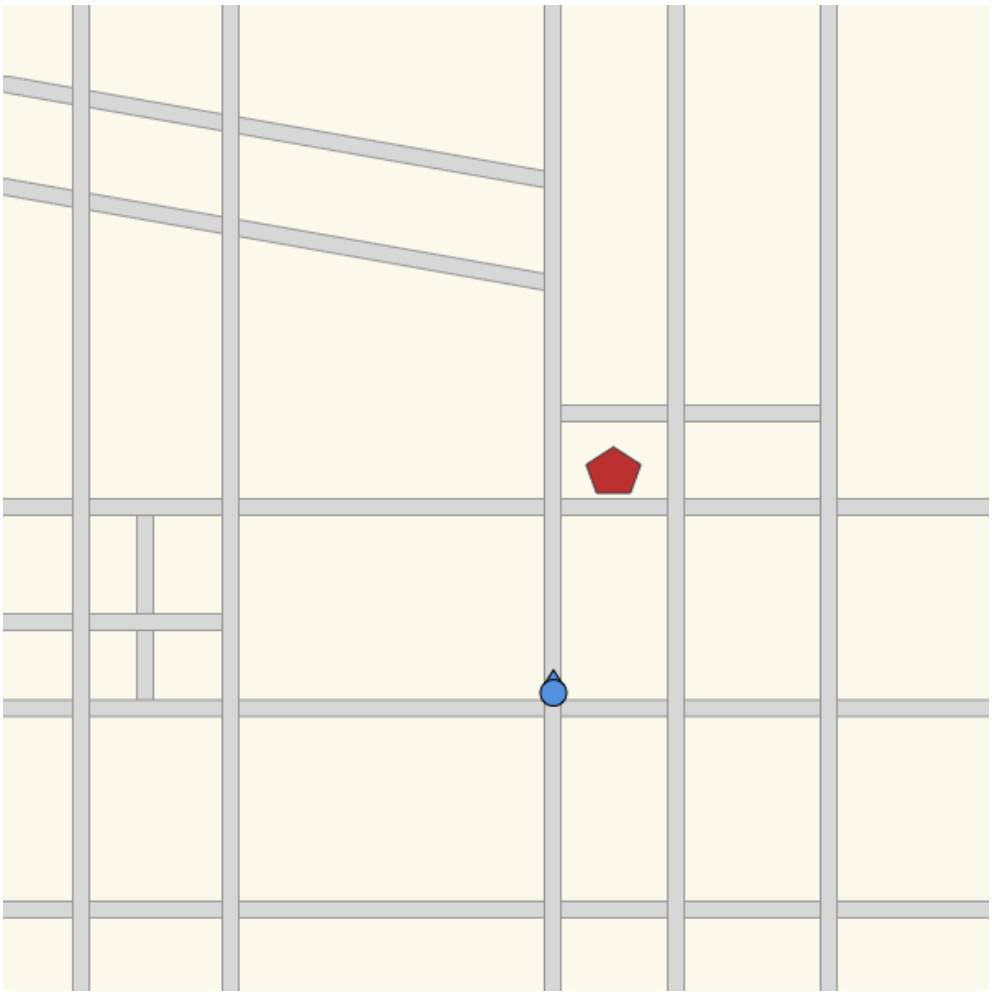

Figure 3 – Example of the GPS navigation animation that participants monitored after selecting a route. The blue circle indicated the current position of the participants vehicle and the red pentagon indicated the destination. The duration of each animation was approximately 15 s and concluded when the vehicle symbol arrived at the destination symbol.

The entire experiment procedure including the decision-making task and the GPS navigation animation were programmed using the Qualtrics web-based survey development platform (Qualtrics, Provo, UT). Participants indicated their preference for a given route by clicking a checkbox next to their desired route display. Responses were

automatically recorded along with the total time each route was viewed until the participants indicated a choice.

2.1.3 Research Design

Each participant completed a total of 16 trials which were divided into four blocks. Blocks were created for the purpose of controlling order effects and there was no discernable difference between blocks to participants. On each trial, participants first viewed a set of four different routes, each described by four attributes (i.e., fuel efficiency, speed, traffic flow, and road quality), then chose one of the four as their preferred route. After choosing a route, participants viewed a GPS map animation which simulated a route navigation monitoring task. After participants arrived at their destination, they received confirmation that the package had been delivered, thus concluding the trial.

The experiment used a 2 (sequence; descending order of alternatives, randomized order) by 2 (salience; highlight attribute label and value, no highlighting) repeated measures design. The dependent measure was the participants choice of route. To examine the effects of sequence and salience on choice, participants' route choice data were recoded into a binary variable based on whether participants chose the route that was highest on the targeted attribute versus one of the other three routes in the route set.

Each participant completed a total of 16 trials (one route choice per trial) that were divided into four blocks. Within each block, the order of trials was randomized for each participant. All sequence and salience conditions were equally represented within each block. Table 1 includes a summary of the trial composition, the sequence and salience

conditions, and the corresponding targeted attributes. Participants completed four trials in each sequence and salience condition, and all four attributes were equally targeted.

Within a block, there were 16 unique route displays; however, each block included the same route displays with a few exceptions. First, each route set was equally targeted by the four sequence and salience conditions. For example, in block 1, participants might view Route Set A displayed with the sequence manipulation present; however, in block 2, Route Set A would be displayed with the sequence randomized. Second, every route had a different name each time it was presented. Thus, in block 1, Route A might be named "Jefferson Avenue" whereas in block 2, that same route would have a different name, such as "Spring Street". Routes were given unique names for each trial so that participants were less likely to feel as though they were seeing the same routes across the experiment.

Table 1 – Trial Composition for a Given Participant in Experiment 1.

Block	Trial	Route Set	Sequence	Salience	Targeted Attribute
1	1	A	Descending	Highlight	Speed
1	2	B	Descending	None	Fuel Efficiency
1	3	C	Randomized	Highlight	Traffic Flow
1	4	D	Randomized	None	Road Quality
2	5	A	Descending	None	Speed
2	6	B	Randomized	Highlight	Fuel Efficiency
2	7	C	Randomized	None	Traffic Flow
2	8	D	Descending	Highlight	Road Quality
3	9	A	Randomized	Highlight	Speed
3	10	B	Randomized	None	Fuel Efficiency
3	11	C	Descending	Highlight	Traffic Flow
3	12	D	Descending	None	Road Quality
4	13	A	Randomized	None	Speed
4	14	B	Descending	Highlight	Fuel Efficiency
4	15	C	Descending	None	Traffic Flow
4	16	D	Randomized	Highlight	Road Quality

Four versions of the experiment were created to address potential carry-over effects associated with participants viewing the same route set four times across the experiment. The only difference between each version was the order of blocks; this allowed for an assessment of the sequence and salience conditions across versions for the first block of trials. This ensured that all four sequence and salience conditions were equally represented for each route set.

Sequence was manipulated to examine how the order of processing alternatives can impact the choice share of a targeted alternative. The sequence manipulation determined the precise order in which participants viewed the four route alternatives. Specifically,

sequence was manipulated by ordering alternatives by descending value of an attribute, or in a randomized order. For example, routes could be sequenced by the fuel efficiency attribute in descending order; consequently, the route with the highest fuel efficiency value would be displayed first in the sequence. In this example, the sequence factor is targeting the fuel efficiency attribute (i.e., fuel efficiency would be the referred to as the targeted attribute).

The salience factor was operationalized as highlighting the attribute label text and the value within each route display (see yellow highlighting in Figure 2). When the salience manipulation was present, one of the four attributes constituted the targeted attribute (e.g., fuel efficiency) which meant the corresponding attribute label and its value were highlighted in yellow. For trials with the salience manipulation present, the targeted attribute was highlighted on all four route displays; therefore, if road quality was highlighted on one display, it was also highlighted on the other three route displays for that trial in the exact same manner. In contrast, for trials when the salience manipulation was not present, no part of any route displays in the route set was highlighted.

Sequence and salience conditions were associated with a targeted attribute. That is, the way in which each manipulation was instantiated depended upon which of the four attributes was targeted. Thus, when sequencing the routes by descending fuel efficiency values, fuel efficiency would constitute the targeted attribute. For trials in which both sequence and salience manipulations were present, then both factors targeted the same attribute. For such trials, if the targeted attribute were fuel efficiency, then the four routes

would be sequenced by descending fuel efficiency values and the fuel efficiency attribute label and value would be highlighted on all four routes.

2.1.4 Procedure

After enrolling in the experiment, participants were free to begin the experiment at their convenience by accessing the URL provided by the SONA system. All experiment procedures were presented online; therefore, participants were able to complete the experiment at their preferred time and location using a computer and web browser of their choice.

Each session began with participants reviewing the informed consent form displayed on the first page of the web-based survey. After obtaining informed consent, participants were presented with an overview of the task procedure, detailed definitions of each component of the task, and specific instructions informing them of their goals for the experiment. Furthermore, participants were given precise definitions of the route attributes as well as example images of a trial timeline, route displays, and the GPS route navigation animation. An overview of the task information and instructions as they were presented to participants is available in Appendix B.

Participants were instructed to choose a route based on the four attribute values. Assuming the role of a delivery service agent, participants were instructed to choose the route that they believed would lead them to arrive at their destination and deliver packages in a timely manner. However, they were told they should also choose a route that ensured their vehicle would operate efficiently and not encounter damage to their cargo. In effect,

participants were given the goal of choosing a route that was high on all four attributes to reduce the likelihood that they would choose a route based on one attribute, regardless of how the decision environment was designed.

In the task instructions, participants were informed that on occasion, some route information would be highlighted in yellow to explore how display design factors might help people make decisions. However, they were explicitly told that the highlighted information did not indicate that the information was more or less important. Moreover, in an effort to prevent unintended bias towards the highlighted information, participants were told that their goal would remain the same on each trial, regardless of whether information was highlighted in the route display.

After reviewing the task instructions, participants completed a practice trial in order to familiarize themselves with the task. The practice trial was identical to an experimental trial, but the salience manipulation was absent; that is, no attribute information was highlighted on the practice trial. Upon completion of the practice trial, participants began the experiment proper.

A timeline depicting the sequence of events for a given trial is presented in Figure 4. Each trial could be divided into three main phases: 1) initial route display, 2) choice, and 3) navigation. During the initial route display phase, participants saw each route in the route set, one at a time as each route was presented on a separate page (i.e., organized by alternative). After reviewing the first route, participants pressed the "next" button to view the second route in the sequence. During the initial display phase, participants could not

return to the previous route, they could only press the next button to view the next route. Furthermore, participants could not choose a route during the initial display phase. Choosing a route was restricted to the choice phase. This restriction was implemented to ensure that participants were exposed to all four routes before making their choice. After viewing all four routes on separate pages, participants began the choice phase.

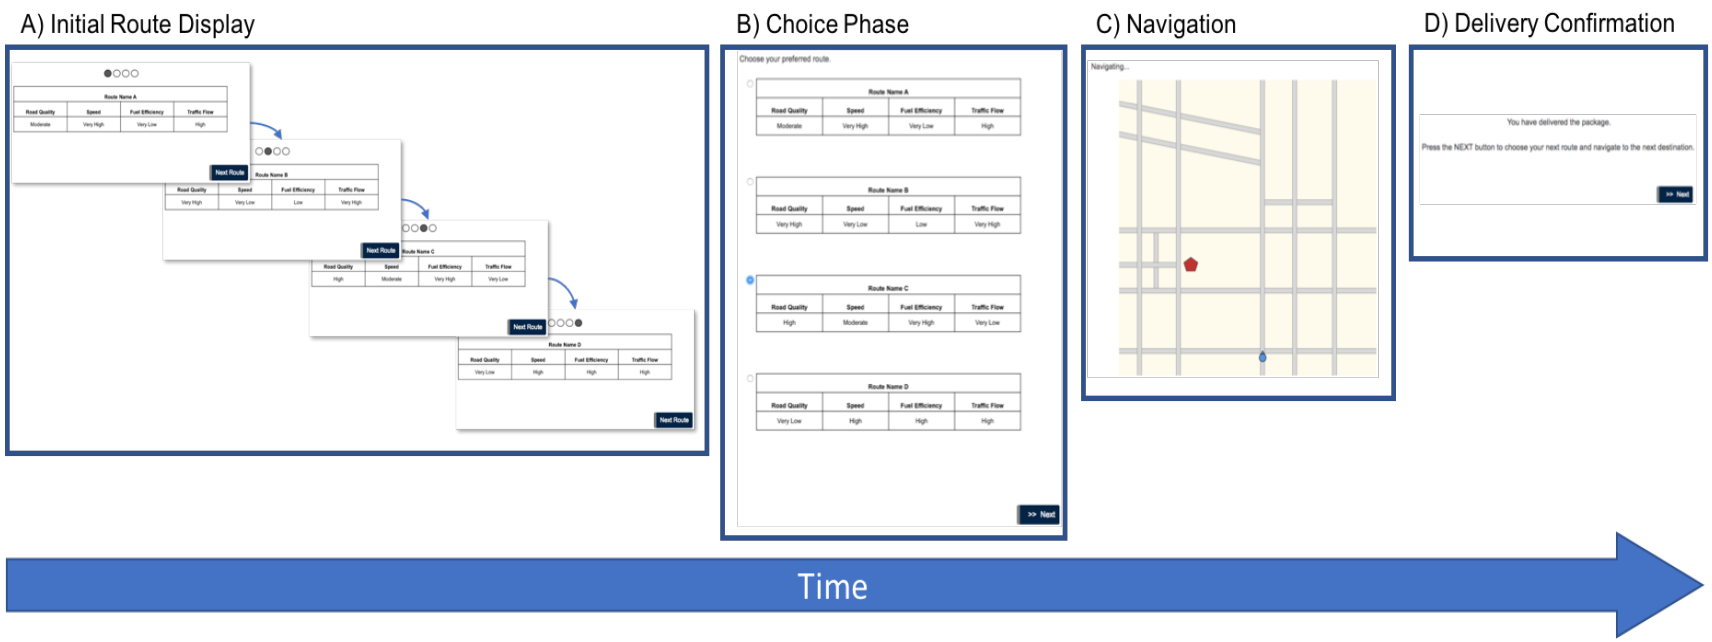

Figure 4 – Trial timeline for experiment 1, depicting the order of events for a given trial. During the initial route display phase, participants viewed each route on separate pages, then pressed the "next" button to view the next route in the sequence. During the choice phase, all four routes were displayed on a single page and participants were able to select their preferred route.

During the choice phase, all routes were presented on a single page and participants could choose a route by clicking a check box next to their preferred route to indicate their choice. If sequence and salience were manipulated in the initial display phase, then they were also manipulated in the choice phase. Thus, the only differences between the initial route display and choice phases was that participants could view all four routes on a single page and they were able to choose their preferred route. Upon choosing a route, participants pressed the "next" button, marking the beginning of the navigation animation phase.

During the route navigation animation phase, participants monitored their vehicle's progress along the route as it navigated to the destination. The navigation animation lasted approximately 15 *s* and the participants' vehicle was depicted as a blue circle with a point indicating the vehicle's direction of travel. After participants' vehicle arrived at the destination marked on the map, a confirmation screen was displayed to indicate that the package had been delivered, thus concluding the trial. At this point, participants were instructed to press the "next" button to choose their next route to their next destination.

All participants viewed 16 identical navigation animations in the same sequence across the experiment, regardless of which route participants chose. In doing so, this ensured equivalence in terms of feedback regarding the participants choice of route; therefore, any interpretation of causality between participants' chosen route and the manner in which the navigation was depicted would be coincidental. To clarify, participants were not given any explicit feedback regarding the accuracy or validity of the route that they chose. Thus, it was unlikely that the animation would constitute feedback that subsequently influenced how participants interpreted route attribute values and their route preferences.

Upon completion of the 16 trials, participants completed a post-study questionnaire. In addition to demographic information, the questionnaire included items assessing participants' perceived importance of each route attribute on a scale from 0 to 100. After completing the questionnaire, participants pressed a submit button, were thanked for their participation, and returned to the SONA website where they were automatically granted SONA credit.

2.2 Results and Discussion

Participants' route choice data were analyzed to determine how sequence and salience influenced choice of routes. Each route choice was coded as a binary variable referred to as targeted route choice, which was based on whether participants chose the route highest on the targeted attribute (i.e., highlighted attribute or sequencing attribute) or one of three route alternatives in the route set. A binary logit regression model was used to examine sequence and salience effects (e.g., Dhar, Nowlis, & Sherman, 2000). The model's dependent variable was whether participants chose the targeted route, as a function of sequence, salience, and the two-way interaction between sequence and salience. The targeted route choice share for each sequence and salience condition is presented in Figure 5.

There was a marginally significant effect of salience on targeted route choice share, χ^2 (1) = 3.41, p = .064. Highlighting a targeted attribute led to a 4% increase in choice share for the route highest on the targeted attribute. Although only marginal, the observed salience effect supports the interpretation that highlighting increases attribute salience and

encourages participants to selectively process the highlighted attributes. Previous research regarding salience effects showed that increasing attribute salience led to a similar effect in terms of encouraging people to selectively process alternatives based on the salient attribute (Jiang & Punj, 2010).

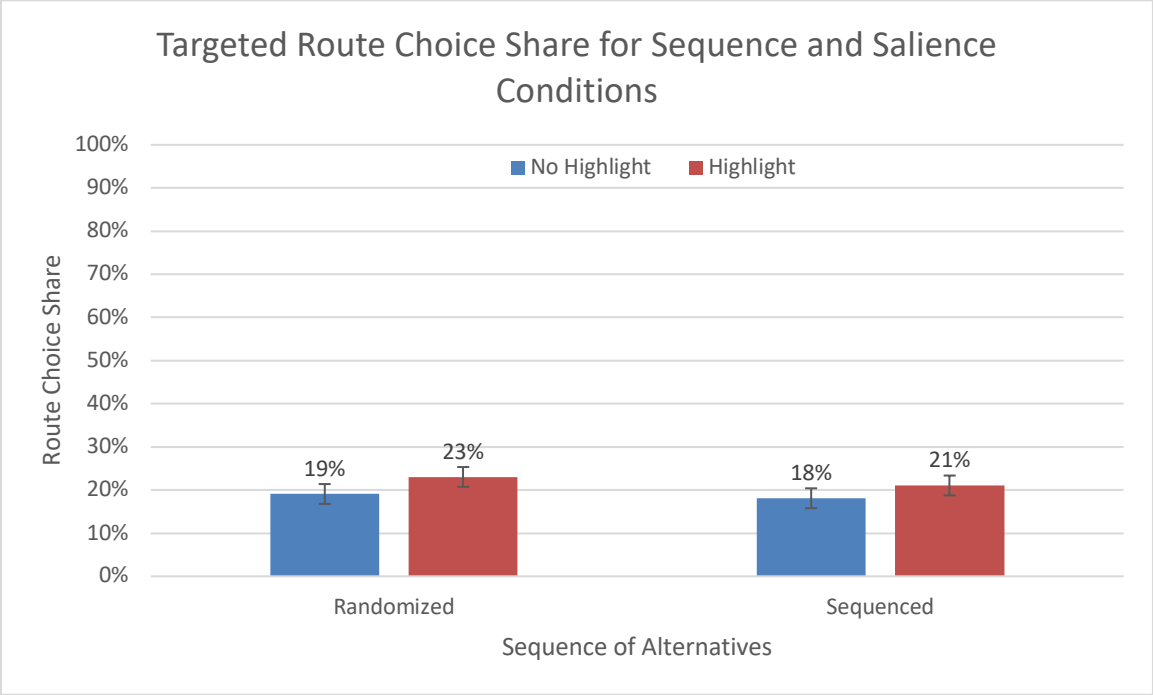

Figure 5 – Targeted route choice share for all sequence and salience conditions, depicting a marginally significant main effect of salience on targeted route choice share. Illustrates how highlighting a targeted attribute can influence people to choose an alternative that is highest on the highlighted attribute.

In addition, the salience effect observed in the current experiment demonstrates that highlighting might be sufficient for influencing choice. Jiang and Punj (2010) used highlighting in combination with increased font size; therefore, the effects of highlighting were confounded with font size. By isolating salience to include only highlighting, the

current experiment provides some evidence that highlighting alone might be sufficient to influence choice.

Analyses revealed there was no sequence effect alone nor in combination with salience on participants' choice of routes. Participants choices were not influenced by whether the routes appeared in a randomized order in comparison to sequenced by descending value on a targeted attribute. Thus, results failed to support the hypothesized main effect of sequence and sequence by salience interaction.

Before ruling out the potential for sequence effects in the current decision-making paradigm, it is important to review how the choice phase procedure was designed. In the current study, the decision-making task was divided into two phases: an initial display phase and a choice phase. During the initial display phase, each route display was presented on a single page and participants could not choose their preferred route. For the choice phase, participants viewed all four routes on a single page with the ability to choose their preferred route by clicking a checkbox next to that route (see Figure 4 for a depiction of the choice phase).

The purpose of separating the task into initial display and choice phases was to ensure that participants viewed all four routes in the order that was prescribed by the sequence manipulation, before they could indicate their choice. However, an unintended consequence of doing so was that participants could have viewed the initial display phase as an unnecessary step in the process, because all routes would be presented on a single page during the choice phase anyways. It is possible that participants ignored the decision

information during the initial display phase and only began truly evaluating the routes during the choice phase. Therefore, it is unclear whether participants experienced the sequence manipulation as it was intended. Experiment 2 further investigated sequence effects to rule out this alternative explanation.

CHAPTER 3. EXPERIMENT 2

Building upon experiment 1, the goal of experiment 2 was to validate the salience effect and to explore the alternative explanation for null sequence effects. To address this alternative explanation, experiment 2 involved changing the choice phase procedure to rule out the explanation that participants did not evaluate the routes in the order prescribed by the sequence manipulation. If a sequence effect were observed in experiment 2, this would indicate that the null effect in experiment 1 was likely due to the fact that participants were free to evaluate the routes in any order they preferred during the choice phase.

With regard to salience, experiment 2 was designed to address two research objectives. First, replicate the salience effect observed in experiment 1. The salience effect on targeted route choice was marginally significant; therefore, validation was necessary to bolster confidence in the salience effect before ultimately applying these effects to the design of an automated DSS. Second, replicate and extend the salience effect to include decision tasks involving stacked displays. There are many applications in which screen real estate is limited (e.g., smartphones and in-vehicle displays). To address such scenarios, designers can stack information displays within the screen, as opposed to distributing them across a single screen (Jang, Trickett, Schunn, & Trafton, 2012). In experiment 2, the modified choice phase presented one route per page, thus creating a decision-making procedure that involved displaying routes in a stacked manner (i.e., one route display presented at a time, in a serial manner). Demonstrating salience effects in a stacked display context would extend its application to additional use cases and display design contexts. If

no salience effect were observed, then it might suggest that salience effects are constrained to applications in which displays are distributed across a single screen.

3.1 Method

A total of 80 undergraduates (42 female, M_{age} = 20.24 years, SD_{age} = 2.96) participated in this experiment. The approximate 30-minute experiment was completed in partial fulfillment of a research familiarization requirement. Experiment 2 used the same method as experiment 1 with the exception being a change to the trial procedure. Specifically, the choice phase procedure was changed to ensure participants were viewing the routes in the same sequence as during the initial route display phase. To this end, the choice phase involved presenting one route per page, just as in the initial display phase. Thus, in experiment 2, the only difference between the initial display and choice phases of the trial was that in the choice phase, participants were able to click a checkbox next to their preferred route whenever it was displayed on their screen. Figure 6 depicts the trial timeline for experiment 2 and illustrates this change in the choice phase of the trial procedure.

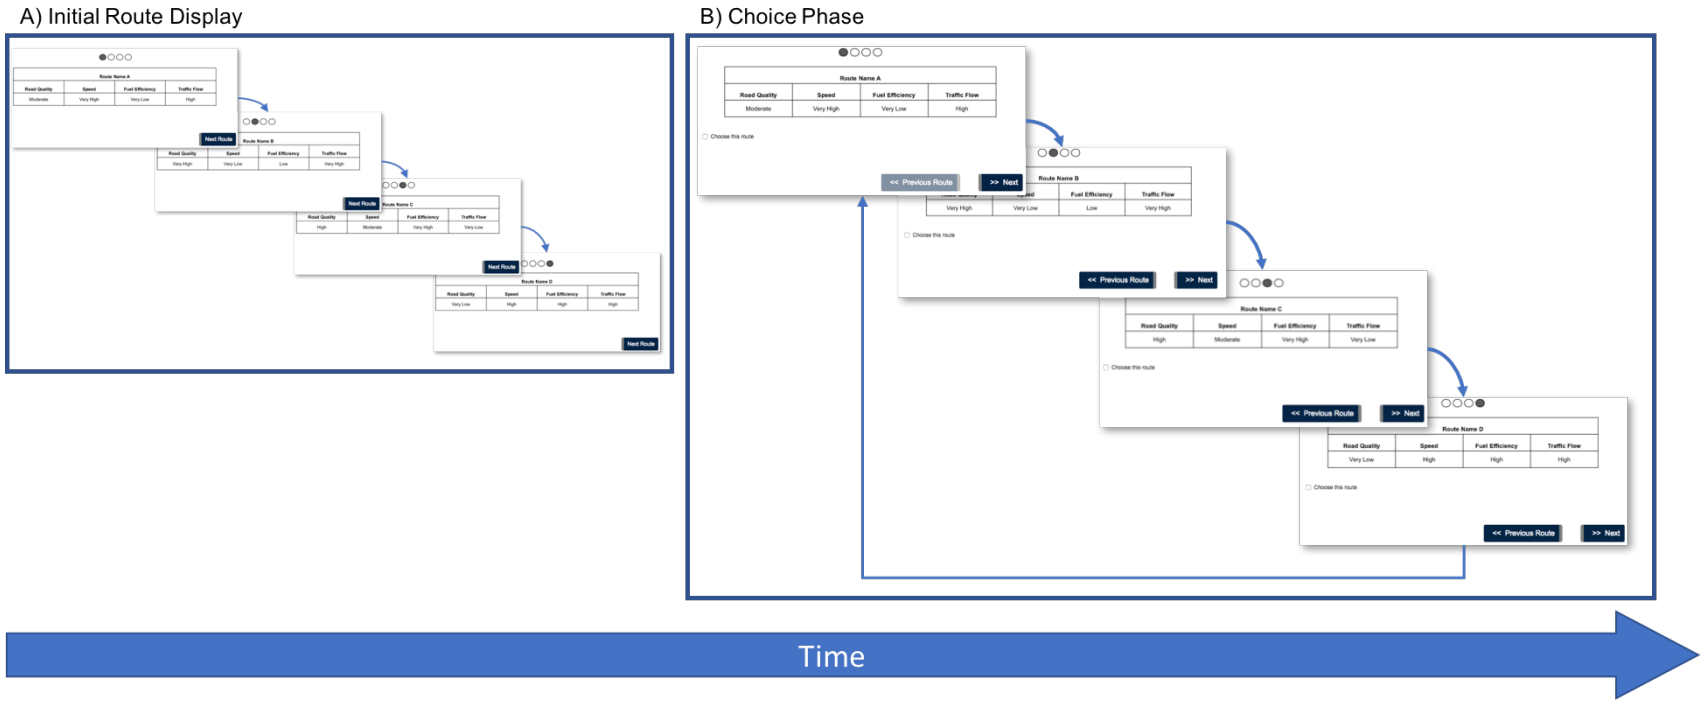

Figure 6 – Initial display and choice phase for experiment 2. Arrows within initial display and choice phases correspond with a required user interaction before changing the page. Illustrates how participants only viewed one route per page during both initial display and choice phases, in contrast to experiment 1 in which all routes were displayed on a single page during the choice phase.

3.2 Results and Discussion

To examine potential sequence and salience effects, participants' route choices were analyzed. As in experiment 1, route choices were coded as a binary variable based on whether participants chose the targeted route or one of the other three alternatives in the given route set. Targeted route choice share was analyzed using the same binary logit regression model as experiment 1 and examined how targeted route choice share varied as a function of sequence, salience, and their interaction. The targeted route choice for each sequence and salience condition is available in Figure 7.

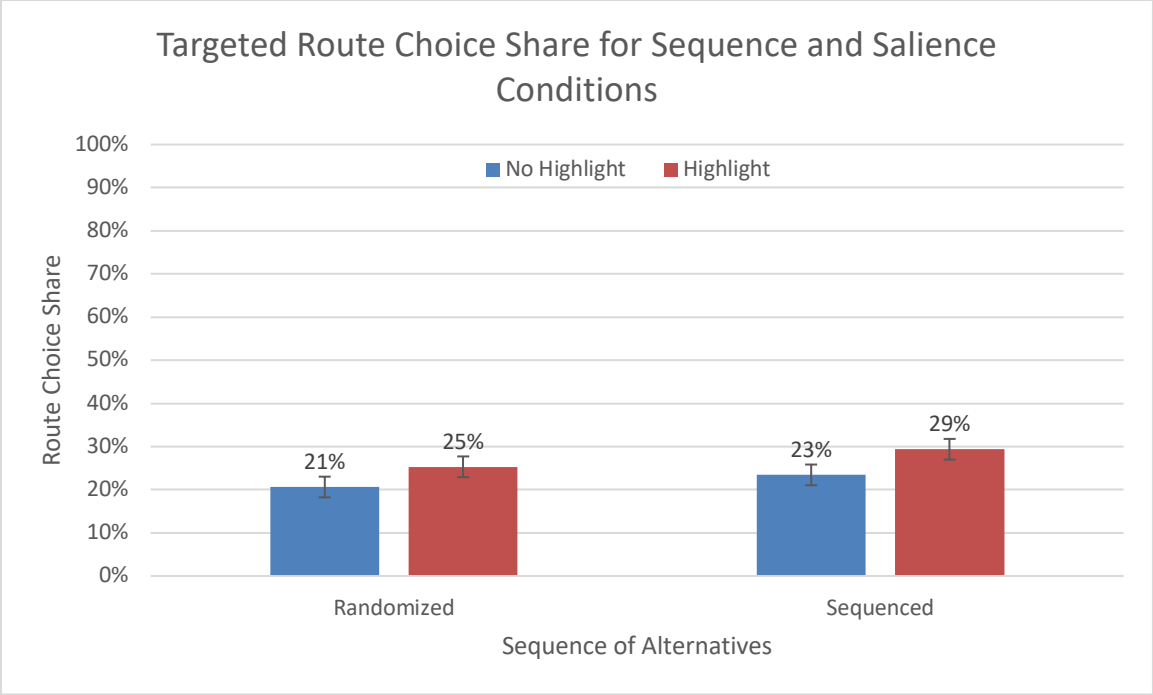

Figure 7 – Targeted route choice share for all sequence and salience conditions. Illustrates how salience increased choice share, regardless of how alternatives were sequenced.

There was a significant effect of salience on targeted route choice share, χ^2 (1) = 5.35, $p < .05$. As depicted in Figure 8, choice share for the route highest on the targeted attribute increased by 5% when the targeted attribute was highlighted. This salience effect validates the marginally significant effect observed in experiment 1 and extends it to include the context of stacked displays. The salience effect can be explained by encouraging selective processing of a targeted attribute, influencing people to choose the alternative that is highest on the targeted attribute when they might not have otherwise. Taken together, these results justify manipulating salience in the design of an automated DSS in order to increase compliance with the automation.

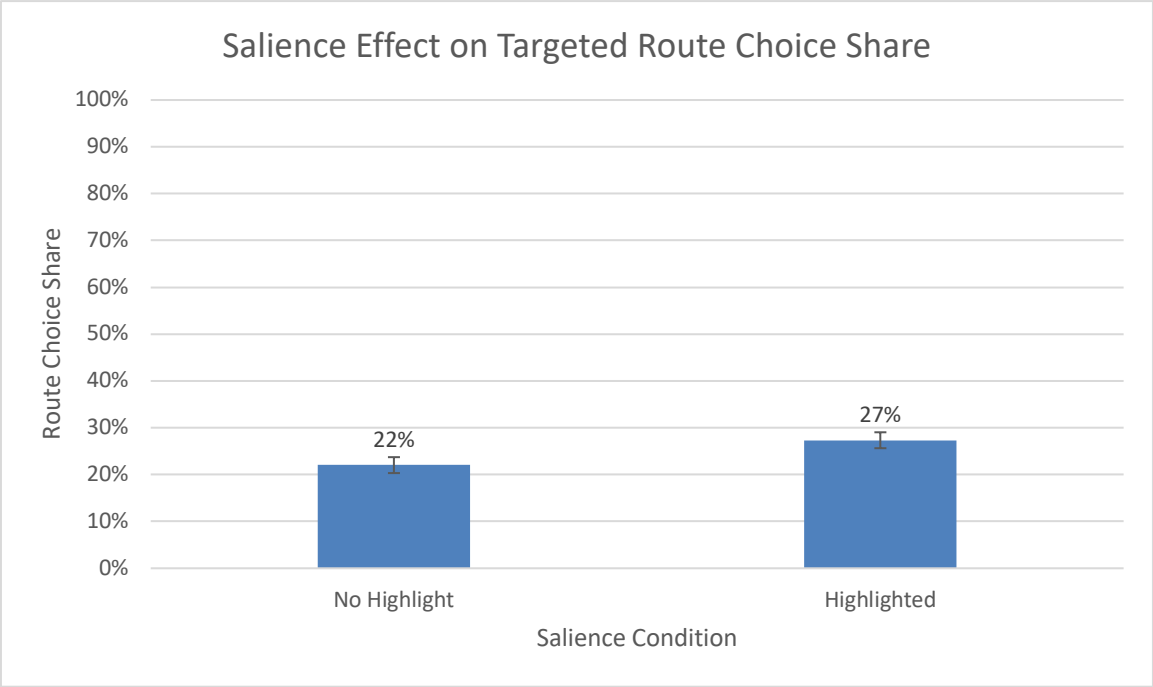

Figure 8 – Main effect of salience on targeted route choice. Highlighting a targeted attribute increased the choice share for the route highest on the targeted attribute. Validates the marginally significant salience effect observed in experiment 1.

Demonstrating salience effects on choice in a decision environment in which information was stacked, as opposed to distributed, bolsters the ecological validity of these salience effects. Previous research on salience effects in multi-attribute choice decision-making tasks have used grouped or list organization containing all the alternatives and their attributes. However, it is not always feasible to make all the decision information readily available to the decision maker within one display; that is, sometimes it must be stacked and require user intervention to change what is displayed. Experiment 2 demonstrated how salience effects can occur with such stacked displays; thus, salience effects can be applied to decision environments that are characterized by reduced screen real estate.

Analyses revealed no significant effects of sequence, neither alone nor in combination with salience. Thus, despite changing the decision-making procedure which forced participants to view each route in the prescribed order, sequencing alternatives by descending targeted attribute value had minimal impact on choice. The lack of sequence effects in experiment 2 indicate that the null effects observed in experiment 1 were not simply explained by the decision-making procedure.

The null sequence by salience interaction suggests that combining these factors does not enhance the relative influence of each factor. However, the results from experiments 1 and 2 showed that combining sequence with salience did not minimize the salience effect. The implication is that if both factors target the same attribute, designers can combine sequence and salience manipulations without ramification because doing so is unlikely to mitigate the relative influence of either factor. Given that designing a decision

environment requires sequencing the alternatives in some manner, this null interaction suggests that sequence can be combined with other factors without interference.

CHAPTER 4. EXPERIMENTS 3 AND 4: INFLUENCING CHOICE IN AUTOMATED DECISION SUPPORT SYSTEMS

The goal of experiments 1 and 2 was to determine how sequence and salience should be manipulated in the visual display used for an automated DSS. To this end, experiment 1 established and experiment 2 validated that highlighting was a viable method of increasing attribute salience to significantly influencing choice. For sequence, results failed to support the notion that manipulating the sequence of alternatives would significantly influence choice in favor of the targeted alternative. Although sequence had no effect on choice, it is worth noting that sequence did not mitigate or minimize the salience effects when both factors targeted the same attribute. This was evidenced by the lack of a sequence by salience interaction in both experiments 1 and 2. Armed with an understanding of how to design a decision environment that predictably influences choice in favor of a specific alternative, experiments 3 and 4 investigated how salience effects can be applied to influence choice and increase compliance with an automated DSS.

Research has shown that introducing an automated DSS can improve decision-making performance, but imperfect automation can also introduce deleterious effects in the form of automation bias (Goddard et al., 2011). To address issues of automation bias, researchers have shown that displaying additional information to increase automation transparency can help reduce automation bias and encourage appropriate use of automation (e.g., Ososky et al., 2014; Rovira et al., 2014). Thus, by helping operators understand why an automated DSS is recommending one alternative over another, it is tenable that

manipulating informationally equivalent display design factors can increase automation transparency.

Chen et al. (2014) argued that increasing automation transparency necessarily requires displaying additional information. However, displaying additional information can lead to increased workload and time spent processing the information (e.g., Dorneich et al., 2017; Wright et al., 2016b), as well as require greater effort from automation designers to implement. For this reason, it would be beneficial to identify informationally equivalent display design factors that could influence choice and increase automation transparency. As experiments 1 and 2 demonstrated, manipulating informationally equivalent display factors like salience can influence choice; however, it is unclear whether salience effects can be used to increase compliance and ultimately increase automation transparency.

Experiments 3 and 4 set out to demonstrate how salience effects could be applied to increase compliance with an automated DSS. Furthermore, by examining salience effects on choice and automation bias, experiments 3 and 4 explored whether salience effects can be used to ultimately increase automation transparency. If so, then manipulating salience would constitute a subtle method of influencing choice and improving how people interact with highly automated systems. In addition, if salience contributes to reducing automation bias, then this would demonstrate how designers can increase transparency without concern for potential side effects associated with presenting additional information.

CHAPTER 5. EXPERIMENT 3

The goal of experiment 3 was to explore how salience effects could be applied to the design of an automated DSS in order to influence choice and increase automation transparency. Results from experiments 1 and 2 were used to inform the display design used in experiment 3. Furthermore, experiment 3 used the same decision-making task as that described in experiment 2 with the one major exception being the introduction of the automated DSS that recommended one alternative on each trial. In experiment 3, sequence was not manipulated as an experimental factor; instead, focus was given to demonstrating how salience effects could be used to increase compliance with the automation's recommended alternative. It was predicted that salience would reduce automation bias. That is, salience would increase compliance with the DSS when its recommended route was superior than others but not when the recommended route was inferior.

5.1 Method

5.1.1 Participants

A total of 80 undergraduates from the Georgia Institute of Technology participated in experiment 3 to satisfy a research familiarization requirement. Participants were recruited by electing to participate using the SONA online experiment enrollment system.

5.1.2 Materials

The stimuli included 48 computer-generated tables, each of which constituted one route display. The design of the route displays was identical to that of experiments 1 and 2. All routes were described by the four attributes used in experiments 1 and 2 (fuel efficiency, speed, traffic flow, and road quality). As in the previous experiments, the four attributes were organized in alphabetical order, from left to right across the route display. In experiment 3, the main difference in materials was that route attribute values were systematically varied to create three different trial types, which is discussed in detail below.

5.1.3 Research Design

The research design constituted a 2 (salience; highlight, no highlight) by 3 (trial type; inferior, equal, superior) repeated measures design. Salience, format, and organization were manipulated in the same manner as in experiment 2. Instead of manipulating the sequence of alternatives, route displays were always sequenced in descending order of the targeted attribute. The dependent measure was participants' route choice which was recoded into a binary measure of whether participants chose the route recommended by the automation or not.

Participants completed a total of 16 trials which were divided into two blocks of eight trials. Blocks were created for the purpose of the experiment design and they were not distinguishable to participants. On each trial, participants viewed a set of four different routes, each of which was displayed on a separate page. After participants viewed each route, they were able to choose their preferred route. This initial display and choice phase followed the same one used in experiment 2.

The main difference in the design was the introduction of the automated DSS. To simulate the DSS, on each trial, the automation recommended one of four routes to participants. The recommended route was always the first route that was presented to participants. For trials in which the salience manipulation was present, the salient attribute always supported the recommended alternative. For example, if the automation recommended Route A and the salient attribute was fuel efficiency, then Route A had the highest fuel efficiency value.

The automated DSS was designed to be an imperfect automated system in order to examine how salience effects might impact rate of compliance and the prevalence of automation bias. To this end, three trial types were created based on the recommended route's attribute utility value relative to the other alternatives in the route set: Equal, superior, and inferior. For equal trial type, each route was equal in terms of overall attribute utility value to ensure that no route constituted a dominant alternative. For the superior trial type, the automation's recommended route had an overall attribute utility value which was superior to the three alternative routes. For inferior trial type, the recommended route had an overall attribute utility value which was lower than all route alternatives in its route set.

As in experiments 1 and 2, a route's utility value was computed by taking the sum of its four attribute values which ranged on a 5-point scale. Therefore, the highest overall attribute utility value a route could have would be 20 (i.e., "Very High" for each attribute) and the lowest would be four (i.e., "Very Low" for each attribute). In experiment 3, the lowest overall attribute utility value was 10 and the highest was 13. A summary of overall attribute utility values for each trial type is presented in Table 2. For superior trial types,

the recommended route was only 1-point higher in terms of overall attribute utility value than the other three routes in its route set. Similarly, for inferior trial types, the recommended route was always 1-point lower in utility value than the alternative routes. Therefore, this created a context in which the recommended route was either equal, slightly better (i.e., superior), or slightly worse (i.e., inferior) than the other routes in the route set. A comprehensive list of the attribute values and overall utility value for each route is available in Appendix F.

Table 2 – Overall Attribute Utility Values for Each Trial Type

	Utility Value	
Trial Type	Recommended Route	Other Routes in Route Set
Superior	13	12
Equal	10, 11, 12, 13	10, 11, 12, 13
Inferior	11	12

Within a block, there were 32 different route displays presented to participants (eight route sets, each of which contained four different routes). Within a session, each participant viewed each route twice, but only once per block. Thus, if a participant viewed route set A with attribute Z highlighted in block 1, then they would view route set A with no attribute highlighted in block 2.

Each block contained four equal, two superior, and two inferior trial types. Trial type was inherent to its route set; thus, it did not vary from block to block. The justification for this composition of trial types within each block was that it allowed for multiple opportunities to measure when the automation made a “mistake” (i.e., automated DSS

recommended an inferior alternative). Moreover, by including twice as many equal as inferior trial types, the intention was to create the perception that the automation was moderately reliable. That is, for equal trial types, no route was clearly better than another; therefore, participants would be able to justify the automation's recommendations on equal trial types, without the need to increase the number of superior trial types. Moreover, influencing choice in favor of a recommended alternative when each alternative is equal in terms of attribute utility provides a stronger test of the salience effects than when the recommended alternative is superior to the other routes.

For trials in which the salience manipulation was present, the highlighted attribute always supported the recommended route. In other words, regardless of trial type, the highlighted attribute (i.e., targeted attribute) was always high. For example, on an inferior trial type, the recommended alternative would have the lowest overall attribute utility value; however, the recommended route's targeted attribute would be the strongest of any alternative in the route set. The reason for this was to prevent the scenario in which the recommended alternative was displayed with its lowest attribute highlighted. Otherwise, doing so could have led to unintended misuse of the salience manipulation.

Four versions of the experiment were created to ensure that each of the four attributes constituted the targeted attribute and that each was targeted on an equal number of trials. Furthermore, each version of the experiment included a different order of the blocks to account for potential carry-over effects. Participants were assigned to one version of the experiment using block randomization. The order of trials within each block was randomized for each participant.

5.1.4 Procedure

All experiment procedures were presented online using the Qualtrics web-based survey platform (Qualtrics, Provo, UT). Experiment 3 followed the same overall sequence of events as the previous experiments and involved the same decision-making task and procedure that was used in experiment 2 with a few exceptions. First, the description of the decision-making paradigm as it was provided to participants was altered to include details regarding an automated DSS which recommended one of the four routes within each trial. Information regarding the DSS was provided to participants at the beginning of the experiment when they reviewed the task instructions.

A copy of the instructions that were provided to participants at the beginning of experiment 3 is available in Appendix C. Participants were informed that an automated system would recommend one route on each trial. Participants were told that the automated system's route recommendations were 70% reliable in terms of recommending the optimal route or equal to other high-performing routes. Participants were informed that the first route in the sequence would always constitute the automation's recommended route and it would be followed by three route alternatives. In addition, text was added to the first route display to communicate to participants that it was the route being recommended on each trial. An example of the recommended route display is presented in Figure 9.

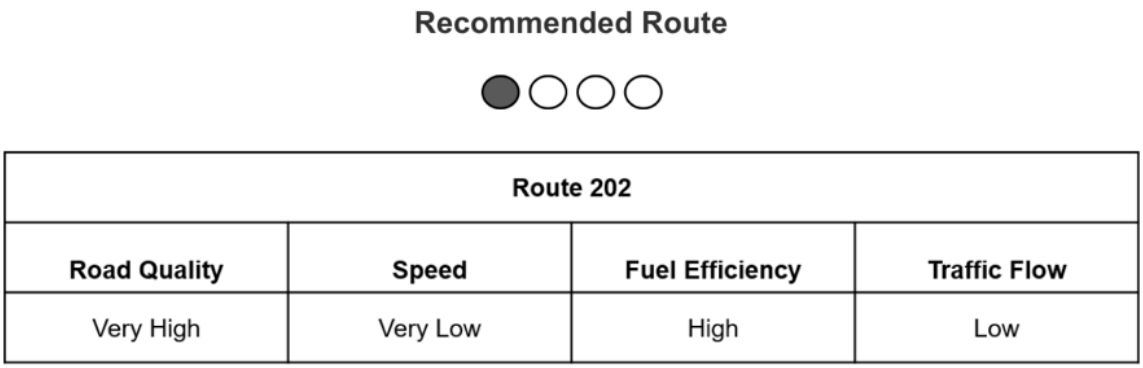

Figure 9 – Example of a recommended route display. The recommended route always appeared first in the sequence on each trial, along with the text "Recommended Route" above the page indicator. Note that this route display was viewed under the "no highlight" condition.

Participants were told to consider the automation's recommended route, but that they should still evaluate all routes and choose the one they believed would lead to the best performance. Participants were informed that a high performing route was characterized as one that ensured the vehicle and cargo were not damaged in transit, that enabled efficient travel to their destination and that the package was delivered in a timely manner. Thus, these task instructions presented participants with the nearly impossible goal of choosing a route that was optimized on all four attributes.

Similar to the previous experiments, participants were given general information about the salience manipulation. Specifically, participants were informed that occasionally, some route's attributes would be highlighted in yellow and that these attributes were chosen by the automated system in order to convey why it was recommending a given route. However, participants were told that highlighting did not convey that an attribute was any more or less important. Moreover, participants were informed that their goal of choosing a

high performing route would remain the same, regardless of whether information was highlighted.

After reviewing the task instructions, participants completed a practice trial in order to familiarize themselves with the task, followed by the experimental trials. The practice trial was identical to the experimental trial, but the salience manipulation was absent (i.e., no attribute information was highlighted on the practice trial). Upon completion of the practice trial, participants began the experiment proper.

Experiment 3 followed the same general decision-making procedure as that of experiment 2 (see Figure 6 for an overview of the trial timeline). In short, each trial was divided into three phases, initial route display, choice, and navigation. Each route was presented on a separate page. During the initial display phase, participants reviewed each route display, then pressed the "next" button to view the next route in the sequence. After reviewing all four routes in the route set, participants began the choice phase, which appeared identical to the initial display phase with the one exception being that participants were now able to choose their preferred route. Participants chose a route by clicking on a checkbox located under each route display, then pressing the "next" button to effectively submit their choice. After choosing a route, participants began the navigation phase of the trial, in which the participant's vehicle was animated to depict it navigating to its destination on a computer-generated map. Experiment 3 used the same animations as experiments 1 and 2. All other aspects of the experiment procedure and the materials was identical to that of experiment 2.

In summary, the only difference in the methods between experiments 2 and 3 was the introduction of the automated DSS and the use of three different trial types. To establish the automated DSS, participants received additional instructions regarding the automation, and the first route display in the route set included text which simply indicated it was the recommended route for that trial (see Figure 9). Three trial types were created by manipulating the recommended route's attribute utility value relative to the others in its route set. Trial type served two main purposes: 1) create a context in which the automation had a moderate level of reliability yet was demonstrably imperfect, and 2) examine whether salience effects could increase compliance without increasing commission errors. Furthermore, examining how salience effects might differ depending on trial type provided a way to examine potential boundary conditions for the salience effect (e.g., if salience effect were limited to one trial type).

5.2 Results and Discussion

Participants' route choices were analyzed to determine how salience influenced whether people complied with the automated DSS. Therefore, route choice data were recoded into a binary variable based on whether participants chose the route recommended by the automation or one of the three route alternatives within that route set. The recommended route choice share data are presented in Figure 10 for each trial type with separate data series for each salience condition.

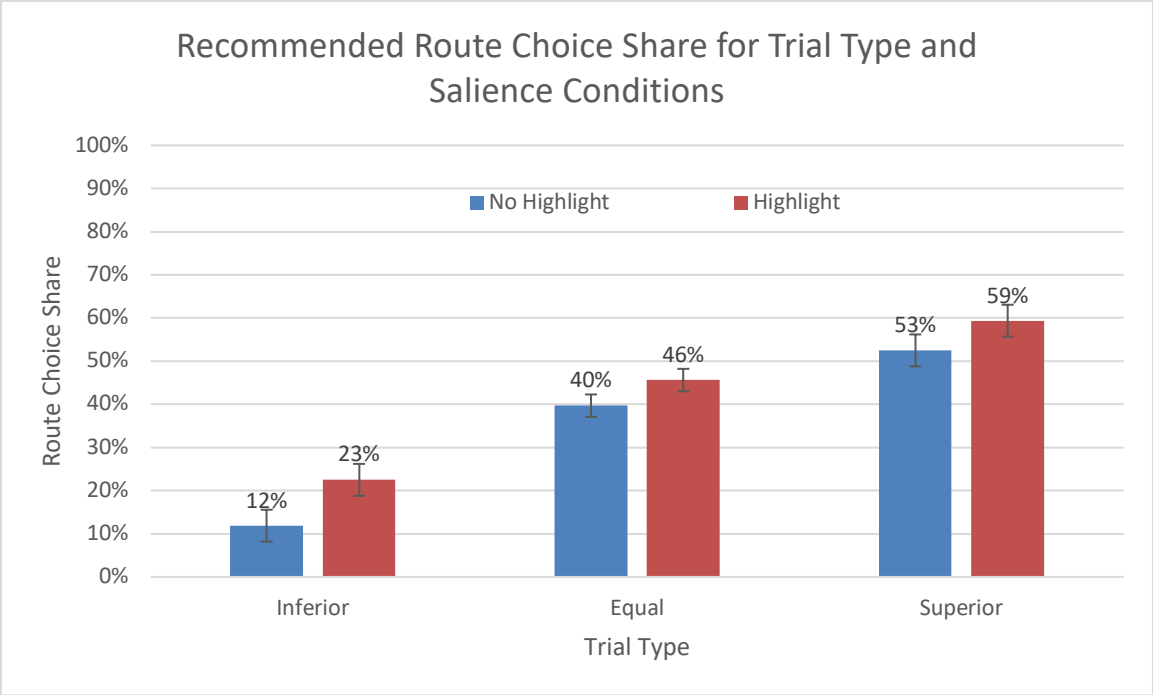

Figure 10 – Recommended route choice share for each salience and trial type condition. Illustrates how recommended route choice share increased due to highlighting, regardless of trial type.

The first analysis examined whether people complied with the automated DSS. If the automated DSS had no impact and people chose routes at random, then the recommended route choice share should be approximately 25%. Across salience and trial type, participants chose the recommended route on 40% of trials, which was significantly greater than chance, t (80) = 7.801, $p < .001$. This result indicates that participants relied on the DSS to inform their choice of routes.

There was a significant main effect of salience on recommended route choice share, χ^2 (1) = 9.86, $p < .05$. The choice share for when the recommended route was not highlighted vs. highlighted is presented in Figure 11. Whether participants chose the route

recommended by the automation depended upon whether the targeted attribute was highlighted (43%) versus not highlighted (36%). In other words, highlighting a targeted attribute which supported the automation's recommended route significantly increased compliance with the automated DSS by 7% across trial type.

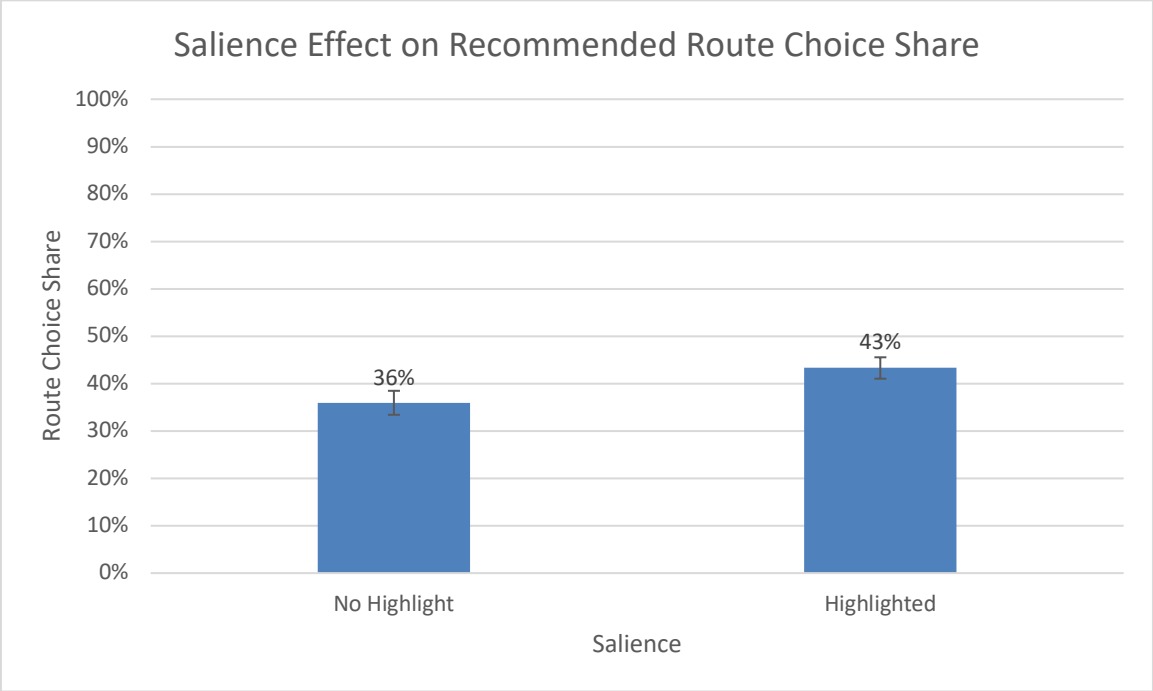

Figure 11 – Main effect of salience on recommended route choice share. Highlighting the targeted attribute, which supported the automation's recommended route, increased recommended route choice share by 7%.

There was a significant main effect of trial type on recommended route choice share, χ^2 (2) = 48.18, $p < .05$. The recommended route choice share for each trial type is displayed in Figure 12. This result illustrates how the recommended route choice share increased as its overall attribute utility value increased. Across salience conditions, recommended route choice share increased from inferior (17%) to equal (43%) and from

equal to superior (56%), *p*s < .01. This trial type effect validates the method of manipulating the recommended route attributes to create a hierarchy based on utility value. Indeed, participants were clearly able to detect the relatively minor changes in overall attribute utility values which were used to create the different trial types. Moreover, the observed trial type effect confirms that participants did not overly rely on the automation.

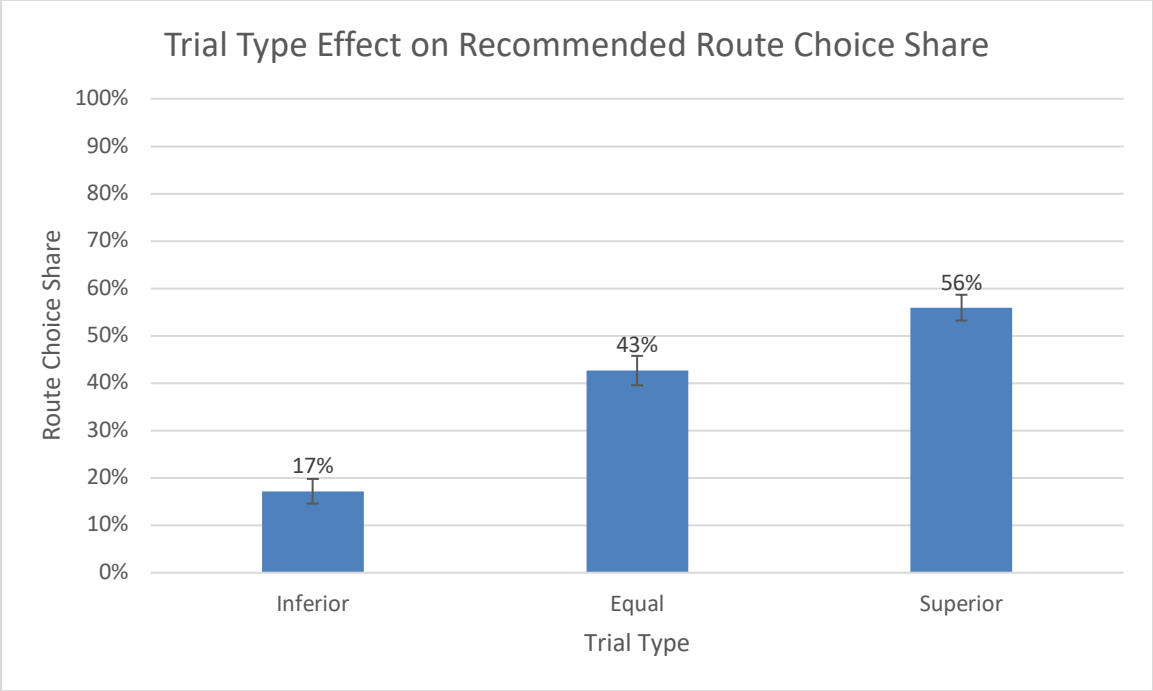

Figure 12 – Main effect of trial type on recommended route choice share. Illustrates how manipulating route attribute utility value to create trial types significantly impacted participants choices.

Results revealed no salience by trial type interaction effect, indicating that the salience effect on recommended route choice share did not depend on trial type. As depicted in Figure 10, salience increased recommended choice share regardless of whether that recommended route was the best alternative available or not. This result suggests that salience increased compliance at the cost of increasing commission errors.

The observed main effect of salience reinforces the explanation established in experiments 1 and 2 which showed that highlighting can be used to increase compliance with an automated DSS. Specifically, increasing attribute salience serves to facilitate selective processing of the salient attribute, increasing the weight given to the salient attribute in the evaluation of each alternative, and subsequently influencing choice in favor of the route highest on the salient attribute. This interpretation can explain how salience effects increase compliance with the automated DSS. In experiment 3, the recommended route was always high on the targeted attribute, even for inferior trial types. Therefore, if salience encouraged selective processing of attributes and increased the weight given to the salient attribute, then it would explain how salience can increase recommended route choice share even for inferior trial types. After all, if participants compared alternatives based on the salient attribute alone, then they would be able to use it as justification for what is otherwise an inferior alternative when considering its overall attribute utility value.

Experiment 3 demonstrated how salience effects could be used to increase compliance but at the cost of potentially increasing commission errors. Highlighting an attribute significantly influenced choice in favor of the alternative recommended by the automation, regardless of whether that alternative was superior or not. Such an effect on compliance is desirable when the alternative is superior; however, increasing compliance only addresses one aspect of automation bias, namely errors of omission (i.e., failure to accept/comply with the automation's recommendation when it is valid). Often, it is equally important to reduce commission errors in order to foster appropriate automation use (Parasuraman & Manzey, 2010).

For designers to leverage salience effects to increase automation transparency, then salience must increase compliance without increasing commission errors. However, results from experiment 3 showed that salience increased compliance even when the recommended route was inferior. Thus, highlighting an attribute might be insufficient for increasing automation transparency. Nevertheless, for applications in which the cost of commission errors is low and increasing compliance is the primary concern, then the results from experiment 3 suggest that salience effects can be applied to influence people to comply with an automated DSS.

CHAPTER 6. EXPERIMENT 4

Experiment 4 investigated how increasing the reliability of the automated DSS might impact salience effects observed in experiment 3. To this end, experiment 4 was designed to address two main goals. First, replicate and extend the salience effects from experiment 3 to include an automated DSS with higher reliability. Given that increasing automation reliability generally increases compliance (e.g., Bagheri & Jamieson, 2004), it is tenable that increasing reliability will serve to reinforce the salience effect and lead to further increases in compliance. On the other hand, increasing reliability might naturally increase compliance leaving little room for salience to influence choice and increase compliance. If such a ceiling effect were to occur, then higher automation reliability would represent an important boundary condition for which salience effects could be leveraged in automation design. Furthermore, automation reliability was increased for the sake of bolstering ecological validity. Considering that automated DSS vary in their level of reliability depending upon the given application, it is necessary to demonstrate that salience effects can be applied to influence choice in systems with increased reliability.

The second main goal of experiment 4 was to explore how increasing automation reliability might impact automation bias. Results from experiment 3 revealed that salience increased compliance with the automated DSS, even when the recommended alternative was inferior. This finding was taken to suggest that highlighting an attribute might be a viable method of increasing compliance but insufficient for increasing automation transparency. However, by increasing automation reliability, it is conceivable that

compliance will naturally increase to the point that salience effects are limited to cases in which the automated DSS recommends a superior alternative. Experiment 4 further examined whether increasing attribute salience can be used as a method to increase automation transparency. To this end, analysis of the salience by trial type interaction is of particular interest. If salience increases compliance for superior but not inferior trial types, then it would indicate that salience can be used to increase automation transparency and reduce automation bias. Alternatively, if it does not, then it would suggest that application of salience effects might be limited to cases in which automation designers wish to increase compliance at all costs, regardless of whether it increases commission errors.

6.1 Method

A total of 83 undergraduates from the Georgia Institute of Technology participated in experiment 4. Participants elected to participate by selecting the experiment from the SONA experiment schedule system. Participants completed the approximate 30-minute experiment in partial fulfillment of a research familiarization requirement.

Experiment 4 used the same research design and procedures of experiment 3 with two exceptions. First, the stated reliability of the automated DSS was increased to 90%. Participants were informed of the automation's reliability during the introduction and overview of the task. Second, to influence participants perception of automation reliability, the number of superior trials was increased from four to eight. Experiment 4 used the same method of manipulating overall attribute utility value to form three trial types as used in experiment 3 (see Table 2 for a summary of the attribute utility values for each trial type).

A summary of the trial type, salience conditions, and targeted attribute which comprised experiment 4 is presented in Table 3 The order of trials was arranged to ensure that the first four trials were always of the superior trial type. Trials 1 through 4 were randomized, but they were always the first four trials presented to participants. The remaining 12 trials were presented to participants in a randomized order. All remaining aspects of the method and task procedure were identical to that of experiment 3.

Table 3 – Example of Trial Composition for Experiment 4

Block	Trial	Route Set	Salience	Targeted Attribute	Trial Type
1	1	A	None	Road Quality	Superior
1	2	B	Highlight	Speed	Superior
1	3	C	None	Fuel Efficiency	Superior
1	4	E	Highlight	Traffic Flow	Superior
2	5	F	None	Road Quality	Superior
2	6	G	None	Speed	Superior
2	7	F	Highlight	Road Quality	Superior
2	8	G	Highlight	Speed	Superior
2	9	H	None	Fuel Efficiency	Equal
2	10	I	None	Traffic Flow	Equal
2	11	H	Highlight	Fuel Efficiency	Equal
2	12	I	Highlight	Traffic Flow	Equal
2	13	J	None	Road Quality	Inferior
2	14	K	None	Speed	Inferior
2	15	J	Highlight	Road Quality	Inferior
2	16	K	Highlight	Speed	Inferior

The purpose of inflating perceived reliability was to examine how salience effects on choice might be impacted by increasing automation reliability. The actual reliability of the automation as participants experienced was about 75%. However, the perceived reliability was inflated by ensuring that the first four trial types that participants

experienced were superior trial types (see "block 1" in Table 3). Participants were told that the system was 90% reliable and the first four trials reinforced this notion that the automation was highly reliable. The justification for separating perceived from actual reliability was that it was necessary in order to collect sufficient data on each trial type, particularly when the automation was "incorrect" as was the case for inferior trial types. Otherwise, to create a system that was truly 90% reliable would require significantly more trials and time to acquire the necessary number of data points.

6.2 Results and Discussion

The recommended route choice share for each trial type is depicted in Figure 13 with separate data series for each salience condition. Across trial type and salience, participants chose the recommended route on 42% of trials, which was significantly higher than chance, t (82) = 7.998, $p < .01$. Although 42% is higher than the 25% that would be expected if participants randomly chose a route, it is lower than one might expect given the stated reliability of 90%. Thus, the observed recommended route choice share indicates that participants did not over rely on the automation's recommendations. In fact, participants frequently disregarded the automation's recommendations. To be clear, this is not to say that the recommendation had no impact on which route participants chose; instead, it is an indication that participants were engaged in the task and appear to have taken their goal of choosing an optimal route seriously.

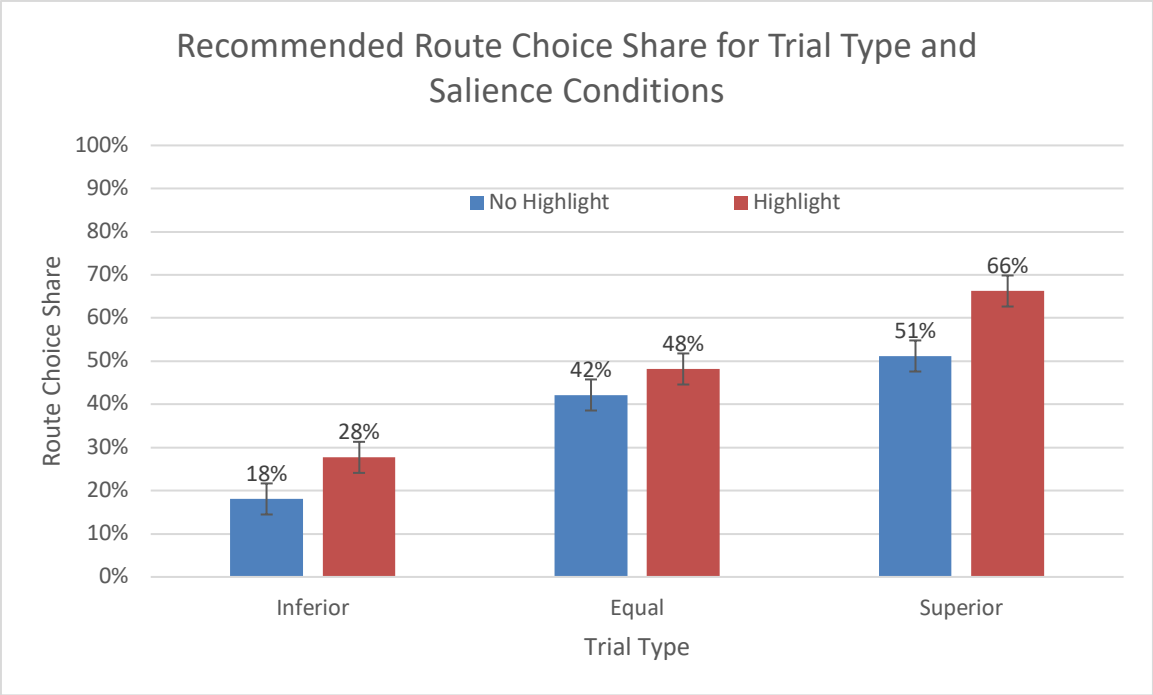

Figure 13 – Recommended route choice share for each trial type and salience condition in experiment 4. When a targeted attribute was highlighted, the recommended route choice share increased across all levels of trial type.

There was a main effect of salience on recommended route choice share, χ^2 (1) = 14.03, $p < .05$. The recommended route choice share for when the targeted attribute was highlighted and not highlighted is presented in Figure 14. When the targeted attribute was highlighted, the recommended route choice share increased by 10%. This result replicates and extends the salience effect of increasing compliance that was observed in experiment 3 to include an automated DSS with increased reliability.

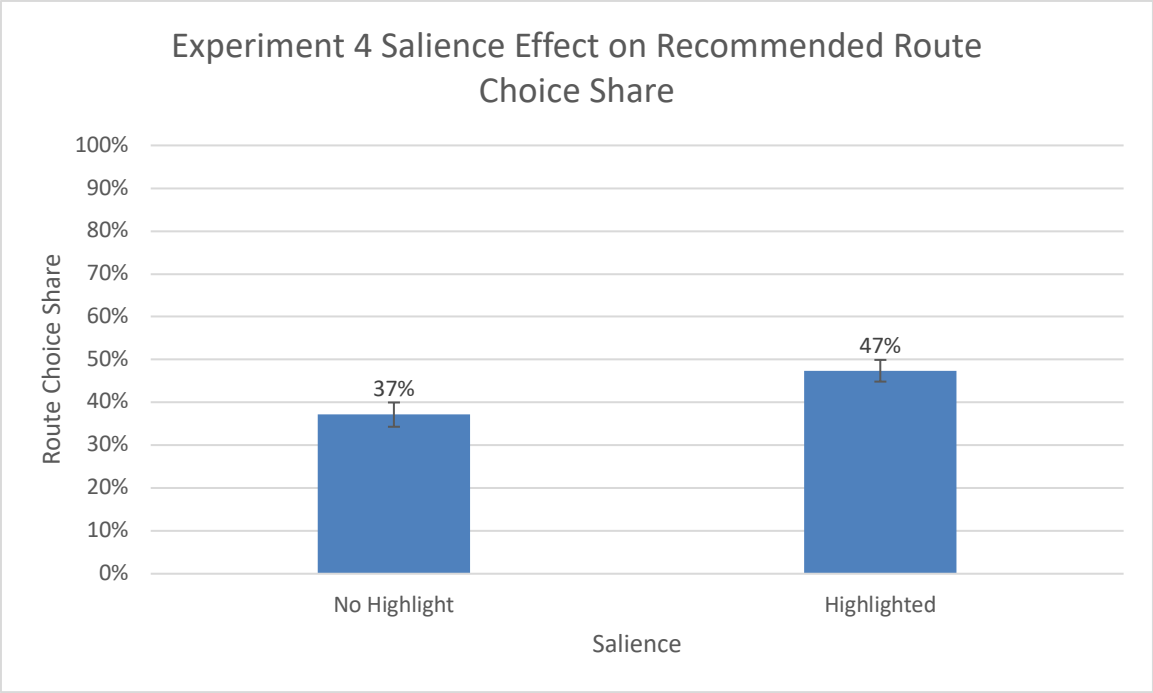

Figure 14 – Main effect of salience on recommended route choice share. Highlighting the targeted attribute that supported the automation's recommended route increased the likelihood that people would choose that route.

The recommended route choice for each trial type is presented in Figure 15. There was a main effect of trial type on recommended route choice share, χ^2 (2) = 37.36, $p < .05$. As expected, choice share increased as a function of the trial type's support for the recommended route. Choice share increased from 23% to 45% to 59% for inferior, equal, and superior trial types, respectively. This result replicates the trial type effect from experiment 3, further validating that manipulating the overall attribute utility value for the recommended route to form distinct trial types had the intended effect. Furthermore, the observed trial type effect indicates that participants evaluated the routes and did not blindly accept the automation's recommendation despite its increased reliability.

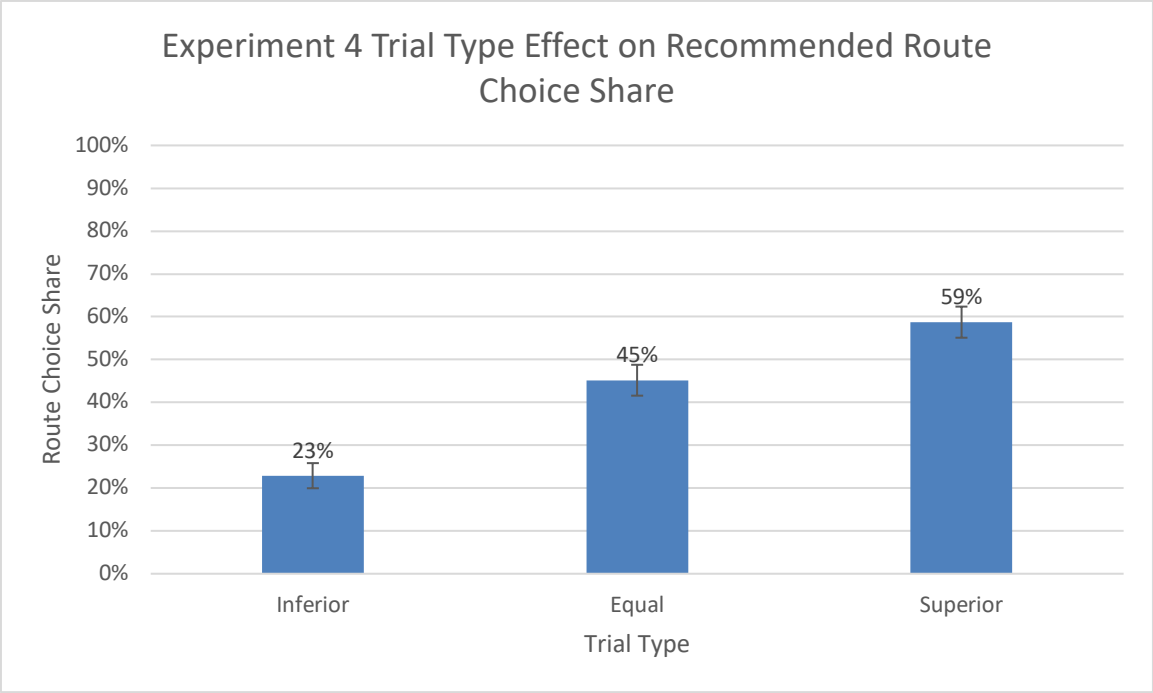

Figure 15 – Main effect of trial type on recommended route choice share. Trial type was defined by the recommended route's attribute utility value relative to other alternatives in the route set. As the recommended route's overall utility value increased, participants were more likely to choose the recommended route.

Analysis of the salience by trial type interaction revealed no significant effect on recommended route choice share. As depicted in Figure 13, highlighting the targeted attribute consistently increased the recommended route choice share across all three trial types. Thus, salience increased compliance with the automation's recommendation regardless of whether the recommended alternative was superior or not. Despite increasing automation reliability, salience increased compliance with the automated DSS when the recommended route was superior, but at the cost of increasing compliance on inferior trial types. It was an open question as to whether an informationally equivalent display design factor like salience could be used to increase automation transparency and reduce

automation bias. However, these results suggest that salience effects might be limited to increasing compliance but insufficient for increasing automation transparency.

In summary, experiment 4 validated that salience effects can be used to increase compliance with an automated DSS despite increasing the reliability of the automation. This indicates that designers can leverage salience effects to foster compliance with an automated DSS. However, careful consideration should be given to the cost of increasing commission errors (i.e., choosing the automation's recommendation even though there is a superior alternative available).

CHAPTER 7. GENERAL DISCUSSION

The goal of this dissertation was to demonstrate how the influence of display design factors could be leveraged to increase compliance and foster appropriate use of an automated DSS. To this end, four experiments were conducted with the overarching goal of designing a decision environment that influenced peoples' choices in favor of a targeted alternative that they might not have otherwise chosen. Moreover, the intention was to identify display factors that could be manipulated while retaining information equivalence, without restricting the individual's decision autonomy. Designing such a decision environment required instantiating the format, organization, sequence, and salience. Previous research guided how to instantiate format and organization, and experiments 1 and 2 examined how to manipulate sequence and salience to influence choice in favor of a specific alternative. After validating that salience could be manipulated to reliably influence choice among alternatives, experiments 3 and 4 demonstrated how salience effects could be applied to increase compliance with an automated DSS.

Experiments 1 and 2 established how increasing the salience of an attribute can influence people to choose the alternative that was highest on the salient attribute. Based on previous research, sequence and salience effects on choice were predicted, yet results revealed only salience increased targeted route choice. Research on salience effects in choice decision making literature has shown that highlighting in combination with increased font size can influence choice among alternatives (Jiang & Punj, 2010). Given the observed salience effects in all four experiments, the current dissertation provides

evidence that highlighting (without changes to font size) is sufficient to increase attribute salience and influence choice.

The observed salience effect can be explained in terms of increasing the selective processing of alternatives based on the salient attribute. By increasing the salience of an attribute, participants selectively processed the salient attribute on each route display. As a result, salience effects impacted how participants weighted the salient attribute and led participants to choose the route that was highest on the salient attribute (i.e., recommended route). This explanation of the salience effect is consistent with similar interpretations of salience effects on choice (e.g., Jiang & Punj, 2010).

As an alternative explanation, one could argue that the salience effects were due to a simple demand characteristic. That is, instead of encouraging selective processing of attributes, highlighting the attribute served to signal to participants which route to choose to support their perception of the hypothesis. In effect, salience led participants to take on the good-participant role (Weber & Cook, 1972). Although no data were collected that refute this explanation, there are two points to consider. First, participants were explicitly told that the highlighted information was not an indication of its importance or validity. Thus, participants would have needed to violate this instruction in order for the salience effect to be attributed to a demand characteristic. Second, in an applied context, such a demand characteristic effect contributes to what the designer is aiming to achieve in that it influences the decision-maker's choice in alignment with the designer's intentions. Nevertheless, one way to rule out the demand characteristic explanation would be to conduct post-experiment interviews inquiring about participants awareness of the

hypotheses. Correlating participants' responses in the interview with their choice data would reveal whether participants knowledge of the hypothesized salience effect was associated with choosing the recommended alternative.

Experiments 3 and 4 leveraged salience effects to increase compliance and foster appropriate use of the automated DSS. The results demonstrated that highlighting a targeted attribute can increase compliance with the automation's recommended alternative. However, salience increased compliance at the cost of increasing commission errors (i.e., complying with the automation on inferior trial types). Specifically, in both experiments 3 and 4, salience increased recommended route choice share, regardless of whether the automation's recommendation was the superior alternative or not. The conclusion drawn from these results was that although highlighting can be used to influence choice and increase compliance, it will not necessarily lead to better overall system performance. That is, for applications in which the cost of commission errors is exceptionally high, then the potential increase in compliance due to salience will be outweighed by the high cost of more commission errors.

One of the goals of experiments 3 and 4 was to investigate how manipulating attribute salience might be used to increase automation transparency. Results from experiments 3 and 4 suggest that increasing the salience of strategically selected information is an insufficient method of increasing transparency. Chen and colleagues (2014) argued that increasing automation transparency necessarily requires displaying additional information to help the human understand the automation's actions. Challenging this notion, the current dissertation explored how informationally equivalent display design

factors can be used to increase transparency without displaying additional information or restricting the decision-maker's autonomy.

For the current study to demonstrate that salience effects could be used to increase automation transparency, salience would need to increase compliance without increasing commission errors. To test this, the salience by trial type interaction was examined in experiments 3 and 4. Results would have needed to show that salience increased compliance for superior and equal trial types yet decreased compliance for inferior trial types. However, the results failed to support this outcome as salience increased compliance across all trial types in both experiments. Therefore, salience can be used to increase compliance, but it may not adequately reduce automation bias in terms of commission errors.

This is not to say that salience effects cannot be leveraged to increase transparency. However, it does suggest that for salience manipulations to effectively increase automation transparency, it must clearly communicate some aspect of the automation's underlying logic. Within Chen et al.'s SAT model, this would correspond with level 2 transparency, the goal of which is to support the human's understanding of why the automation is doing what it is doing. In terms of the current study, this implies that the salience factor would need to clearly convey why the automation is recommending that alternative in order to sufficiently increase transparency. Although this was in fact part of the intention for highlighting attributes in the current study, the results suggest that informationally equivalent display design factors alone might not be able to adequately communicate the rationale and logic underlying an automated system.

Future research should explore how informationally equivalent display design factors like salience might be leveraged to support other display manipulations. For example, in an effort to increase transparency, Rovira et al. (2014) used a color-coded highlighting scheme to communicate to participants the level of automation reliability as it changed from trial to trial. Their results suggest that using salience effects in combination with additional information that communicates the automation's status can help people calibrate trust and foster appropriate use of automation. Rovira et al.'s results illustrate how the salience manipulation needs to clearly communicate a fundamental aspect of the logic underlying the automation and not rely on simply capturing selective attention. In summary, enhancing the perceptual salience of strategically selected information might increase compliance with an automated system's recommendations but it is unlikely to decrease commission errors unless it increases automation transparency.

7.1 Salience Effects on Decision-Making Strategies

The manner in which decision information is displayed can encourage people to use sub-optimal decision-making strategies (i.e., heuristics). People tend to use heuristics in order to reduce effort, which can lead to suboptimal decision-making outcomes (Simon, 1990). Extending this to the current study, increasing attribute salience encouraged people to use a noncompensatory heuristic that consistently supported participants' choice of the automation's recommended route, even when that route was inferior.

In the current study, salience fostered the use of a simple noncompensatory heuristic despite being suboptimal for completing the task. Participants were instructed to

choose a route that was ultimately high on all four attributes. However, route attributes were instantiated such that choosing an alternative that aligned with this goal required making trade-offs among attributes. Given these task instructions, participants should have used an equal weight model that constituted a compensatory strategy. Using such an equal weight strategy would have involved acquiring each attribute, summing the values for each alternative, holding those values in memory to evaluate alternatives, then choosing the alternative with the highest utility value. If participants applied this strategy, they should not have complied with the automation whenever it recommended an inferior alternative. Nevertheless, results showed that manipulating salience led participants to do just that. One reason for this is that people avoid expending the effort required to use a compensatory strategy and instead implement a noncompensatory single-cue heuristic like Take the Best (Gigerenzer et al., 1999; Gigerenzer & Gaissmaier, 2011). As a result, participants simplify the process of acquiring and evaluating attributes (i.e., cues) but doing so comes at the potential cost of choosing an inferior alternative.

Salience increased compliance by fostering the use of a noncompensatory heuristic in which the alternatives were evaluated based on a single-attribute (i.e., cue). Highlighting the targeted attribute indicated to participants a path of least resistance, as opposed to the optimal yet effortful compensatory strategy like the equal weigh model. When the salience manipulation was present, participants simply compared routes by their highlighted attribute values and chose the best one. When the recommended route was inferior, this salience heuristic led participants to choose an inferior alternative. However, for superior trials, salience fostered the use of noncompensatory heuristic that happened to lead

participants to choosing the best alternative because the salient attribute always supported choosing the recommended route, regardless of trial type. When the recommended route was superior, then the salience heuristic happened to yield a valid outcome; on the other hand, when it was inferior, then it led to suboptimal choices. In summary, the salience factor wiped out the use of an equal weight model that more optimal based on the task instructions by influencing people to use a simple and easy to implement heuristic.

To explain how salience led people to consistently use a similar noncompensatory heuristic, consider how heuristics can reduce cognitive effort. For Shah and Oppenheimer (2008, 2009), heuristics tend to influence our decisions by simplifying one or both of cue acquisition and evaluation. For example, instead of considering all attributes in their evaluation of alternatives, decision makers might choose to examine only a subset of those attributes. In a similar sense, display design factors like salience can simplify the acquisition of attributes by highlighting those attributes, capturing selective attention, and leading people examine and evaluate fewer attributes. Applying this to the current dissertation, highlighting the targeted attribute offered a noncompensatory single-attribute heuristic which simplified acquisition and evaluation of attributes. Consequently, alternatives were evaluated based only on their salient attribute, which consistently increased compliance with the automation because the recommended alternative was always high on the salient attribute. In summary, display design factors such as salience can be strategically manipulated, serving to signal to participants the path of least resistance and fostering the use of noncompensatory heuristics, which can lead to predictable effects on choice.

7.2 Limitations

The route navigation paradigm was created to simulate the task of selecting a route from a GPS system and monitoring the vehicle's position along the route. The goal was to leverage an existing paradigm in which participants were familiar with making decisions and that they would be motivated to evaluate the alternatives. However, given the low fidelity simulation and the fact that it was a web-based experiment, it is reasonable to assume that the task might have affected participant motivation. Consequently, salience effects might not generalize to tasks in which the stakes are high or in which people are highly motivated to choose the perceived best alternative at all costs. Despite these concerns, analysis of participants' route choice data showed that they did not randomly choose their routes. Instead, it appears that participants were at least moderately engaged in and motivated to complete the task as they were instructed.

The route displays were partially designed to create a canvas on which sequence and salience could be easily manipulated. This meant that format and organization were controlled across the experiment. Thus, additional research is needed to understand how salience effects might depend upon the use of a table format that is organized by alternative. Although, it reasonable to assume that salience effects would apply to additional display design contexts, the scope of the observed effects cannot be determined based only on results from the current dissertation. However, considering the results from experiments 1 and 2, salience effects did not rely upon direct comparisons between alternatives. This suggests that salience effects would apply to a variety of information organizations beyond

that used in the current study. Additional research is needed to explore potential boundary condition for salience effects.

The effect of salience in terms of increasing compliance might be limited to cases in which the difference between alternatives utility value is small. In experiments 3 and 4, superior and inferior trial types were created to understand how salience effects might differ depending upon the validity of the automation's recommended alternative. Specifically, the recommended alternative's utility value was decreased by one value for inferior and increased by one for superior trial types. As the results showed, salience increased compliance with the automated DSS, regardless of trial type. However, it is worth noting that one utility value was the smallest possible difference between alternatives utility value. Therefore, it is unclear whether salience effects only influence choice when the differences between alternatives is minimal. Presumably, there is a range of differences between alternatives that salience can influence choice, but outside of which salience effects are null. For example, if the recommended alternative is dominated on all attributes, then any motivated decision-maker is unlikely to choose the recommended alternative, regardless of any salience manipulation. However, precisely at what point do salience effects begin to diminish is an open question. Future research should investigate the relationship between increasing attribute utility value and the efficacy of salience manipulations.

Across all four experiments, it is worth noting the relatively small size of the salience effect on choice. The salience manipulation led to a 4%, 5%, 7%, and 10% increase in recommended choice share for experiments 1, 2, 3, and 4, respectively. Depending upon

the application, it could be argued that 4% to 10% is practically insignificant. However, it is worth considering two points. First, these effects on choice were made by simply highlighting an attribute in yellow on each route display. This means that salience effects can be easily implemented by designers in a wide array of applications. Furthermore, leveraging salience effects does not require presenting additional information, nor does it restrict the decision maker's ability to choose. Second, in many cases 4% to 10% is not trivial, especially considering how such effects could compound over time and as the application scales up to include a larger portion of the population. Take for example the Google Maps application for smartphones which has over 1 billion active users (Lardinois, 2016). At such a large scale and on repeated exposure, these small effects could be used to influence many people to make better decisions.

7.3 Theoretical Implications

The current dissertation offers evidence that increasing the salience of an attribute in a multi-attribute choice decision-making task can influence choice by encouraging people to selectively process attributes. By capturing selective attention, salience increases the likelihood that people will choose an alternative based on its salient attribute value. Researchers have proposed similar explanations for salience effects (e.g., Jiang & Punj, 2010; Lurie & Mason, 2007; Mukherjee & Srinivasen, 2013), and the current dissertation provides further evidence for this underlying mechanism. Building upon the extant research, the current study showed that salience effects do not rely upon an organization which facilitates direct comparisons between alternatives' salient attributes. This finding

reinforces the notion that salience effects rely upon encouraging people to selectively process attributes.

Combining sequence and salience factors to influence choice in favor of one alternative does not enhance the relative influence of each factor. In experiments 1 and 2, a sequence by salience interaction was predicted but results failed to support this effect. The rationale for this hypothesis was that both factors targeted the same attribute and combining factors should have reinforced the effects of one another. Research has shown that when sequence and salience target different attributes, then combining these factors mitigates the relative influence of each factor and encourages people to choose an alternative which is moderate on both attributes (Jing & Punj, 2010). Given the limited nature of research investigating sequence and salience effects, the current dissertation offers the first evidence that combining sequence and salience factors that target the same attribute does not produce an additive influence on choice among alternatives. Nevertheless, it is important to note that combining sequence and salience did not minimize the salience effect. These results lead to the conclusion that combining factors which target the same attribute might not strengthen their relative influence, but it is unlikely to minimize it either.

In order to increase automation transparency, display design factors must go beyond capturing selective attention in order to increase compliance and reduce commission errors. Researchers have shown that displaying additional information to convey the logic underlying the automation's recommendations can significantly reduce automation bias (e.g., Mercado et al., 2016). The current dissertation explored whether

subtly changing the way that information was displayed, rather than displaying additional information, could be used to similarly increase automation transparency. Results from experiments 3 and 4 showed that increasing salience can increase compliance. However, increasing compliance only addresses one aspect of automation bias; it is also important to reduce commission errors. The implication is that in order to increase compliance and decrease commission errors, one must validate that the display design factors support the human's understanding of the logic underlying the automation's recommended alternatives.

7.4 Practical Impact

The use of highlighting text can be sufficient to increase salience and influence choice. In the current study, salience was manipulated by simply highlighting a targeted attribute label and value in yellow. Previous research has demonstrated salience effects on choice by manipulating the relative brightness of alternatives (e.g., Milosavljevic et al., 2012) or combined highlighting with other factors like increased font size (e.g., Jiang & Punj, 2010). Isolating the salience effect to highlighting offers a method for designers to leverage salience effects without introducing additional information or restricting the decision-makers choice. Furthermore, the salience effect was measured in terms of whether it influenced choice in favor of a specific alternative. The practical impact is that salience can be used to increase the choice share of a targeted alternative and not just a subset of alternatives. For example, suppose it was desirable for a nurse to choose a treatment option that had the quickest rate of recovery. Treatment options could be displayed to the nurse with the rate of recovery attribute highlighted. The current study suggests that doing so

would increase the likelihood that the nurse would select the treatment with the highest rate of recovery rather than one that was middling to high.

Enhancing the salience of a strategically selected attribute can significantly increase compliance with an automated DSS, but it does so at the cost of increasing commission errors. In both experiments 3 and 4, salience increased the recommended route choice share, regardless of trial type. This effect for inferior trial type means that salience influenced people to choose a route that was inferior to others, which is equivalent to a commission error in the automation bias and HAI literature. Therefore, automation designers should carefully consider whether such salience effects should be applied in a DSS. The following are a few conditions which if met could be used to justify the use of such salience effects in an automated DSS: a) the decision-making context and domain establish that the cost of commission errors is relatively low, b) automation reliability and performance is high such that the rate of commission errors is inherently low, and c) the automated DSS is supporting preference-based decision making and performance measurements are subjective or abstract. This latter point warrants further explanation. In preference-based choice, the decision-maker constructs their preferences as they evaluate information, and thus do not necessarily know what option they want to choose (Bettman et al., 1998). In this case, it might be preferable to nudge the decision maker using salience effects and increase compliance, despite the fact an imperfect automation's recommendation will occasionally be inferior to another alternative that is readily available.

For example, consider a farmer operating a combine harvester, which is highly automated but still requires a human-in-the-loop. To run the combine, the farmer chooses a harvesting strategy and the combine automatically adjusts numerous settings in accordance with the harvesting strategy. Choice among harvesting strategies is largely based on preferences and the degree to which the strategy achieves the farmer's higher-level goals for their farming operation. For instance, one harvesting strategy optimizes combine settings adjustments for grain savings and another optimizes for machine efficiency. Farmer preferences will vary but the combine manufacturer might prefer the farmer experience the combine operating at its highest level of efficiency. To this end, designers could present harvesting strategies to the farmer and highlight the attribute which offers the strongest support for the machine efficiency strategy. The machine efficiency strategy might be inferior for a particular farmer on a given day, but the notion of this being a commission error in the broader sense is practically irrelevant.

In summary, there are applications in which the trade-off between increasing compliance by increasing commission errors is less costly or perhaps irrelevant. In such applications, the current study illustrates one way in which salience effects can be applied to influence choice and increase compliance with an automated DSS. Of course, such decisions regarding automation design should not be taken lightly because increased exposure to automation-induced errors can lead to automation misuse and disuse (Lee, 2008).

CHAPTER 8. CONCLUSION

This dissertation presented research at the intersection of display design, decision making, and human-automation interaction to examine how display design can be used to influence peoples' choices and increase compliance. Along the path from automation to autonomy, our interactions with highly automated systems will grow in ubiquity (Hancock, 2017). More and more of our interactions with technology will involve choosing among decision alternatives that constitute plans or general strategies for the highly automated systems to execute. To shape how people might choose among such alternatives, the current dissertation demonstrated how increasing attribute salience can subtly influence peoples' choices and increase compliance with an automated DSS. However, the observed salience effect increased compliance at the cost of increasing commission errors. This suggests that informationally equivalent display design factors like salience might be insufficient for increasing automation transparency, specifically in terms of reducing commission errors. Nevertheless, the current study illustrates how designers can leverage salience effects to subtly yet predictably influence choice in favor of a specific alternative and thus increase compliance with an automated DSS. In doing so, designers can help people to make better decisions and create more enjoyable interactions with highly automate systems, without the human sacrificing their autonomy in the decision process.

APPENDIX A. LITERATURE REVIEW OF INFORMATION DISPLAY FACTORS ON CHOICE DECISION MAKING

In the following subsections, four categories of display factors are defined and research regarding their effects on choice is reviewed in order to inform how to design a visual display to influence choice. In addition, some research regarding the interaction of display factors is reviewed in the final sub-section. Each subsection concludes with a summary of the research and consideration is given to how these display factors can be applied to the design of an automated decision support system in order to influence operators' choices and improve HAI outcomes.

8.1 Format

Decision information can take a variety of forms and the information format category captures this component of the visual display design. For example, decision information regarding different cell phone attributes (e.g., battery life, processor, storage capacity) can be displayed in a table format or bar graph format. Each of these means of describing cell phone attributes represents a distinct form of presenting decision information that can lead to systematic differences in how people evaluate the decision information.

Graphs and tables are examples of two different formats, which research has shown each can lead to differences in terms of how people process decision information (e.g., Meyer, Shamo, & Gopher, 1999). Graphical formats tend to leverage advantages associated

with rapid visual-spatial processing in order to aid decision makers by revealing trends and relationships in the data (Few, 2004). In contrast, tabular formats (e.g., series of rows and columns with separate columns for each attribute and separate rows for each decision alternative) are generally conducive to extracting precise pieces of information. Research has shown that format can influence choice by determining the cognitive effort required to process the decision information (Dilla & Steinbart, 2005; Schkade & Kleinmuntz, 1994) as well as influence the perceived differences between decision alternatives and their attribute values (Sun, Li, & Bonini, 2010).

Dilla and Steinbart (2005) demonstrated a graphical-tabular format effect, illustrating how format can influence choice. Participants were given the role of a business manager tasked with selecting one of four different environmental cleanup plans (i.e., decision alternatives). On each trial, participants were tasked with indicating their preferred cleanup plan based on the range of costs associated with the plan. The format of the four plans was manipulated between subjects as each plan's range of costs were displayed using either a tabular or graphical format. Participants assigned to the graphical format condition chose fewer dominated alternatives (i.e., cleanup plans that were inferior to others in terms of the range of costs) than participants in the tabular format condition. In other words, although there was no optimal cleanup plan, participants in the tabular format condition were more likely to choose a plan that was inferior to others.

Graphical-tabular format effects can be explained by differences in the amount of cognitive effort required to accurately evaluate the alternatives (Dilla & Steinbart, 2005). With a graphical format, people can use rapid, visual-spatial processing to acquire

information about differences between alternatives with minimal cognitive effort. In contrast, a tabular format requires considerable effort to acquire information for one alternative, hold it in memory, and evaluate differences between alternatives. To reduce effort associated with using tables, people used decision-making heuristics that involved integrating less information which does not adequately account for differences between alternatives. Consequently, tables led participants to choose more dominated alternatives in comparison to the graphical format. In short, the graphical-tabular format effect demonstrates how format can determine the level of cognitive effort involved in acquiring information and evaluating alternatives, which in turn influences choice. By leveraging such format effects, designers can increase the likelihood that people will process the information in a particular fashion that leads to the preferred or targeted choice outcome.

To summarize, decision alternatives can be displayed in different forms and systematic differences in the effort required to evaluate information and choose an alternative. Graphs and tables are two instantiations of format that are commonly used to display decision information to people. Graphs tend to influence choice by allowing people to quickly acquire information about fluent differences among alternatives via visual-spatial processing. In contrast, tables tend to influence choice by encouraging people to selectively acquire and process decision information. Due to increased cognitive effort involved in acquiring a comprehensive evaluation of and comparisons between alternatives, tables are more likely to lead people to integrate less information and potentially bias choice. When people seek to conserve effort, tables can enable other factors

to capture attention and influence choice, which might otherwise have no impact when used with graphs.

8.2 Organization

Organization refers to the way in which decision information is arranged within or across displays. Decision information can be organized to form groups or patterns, which can ultimately impact how individuals acquire and evaluate decision information (Kleinmuntz & Schkade, 1993). For example, Russo (1977) used a simple organization manipulation by arranging unit price information associated with consumer products into a list, in contrast with traditional unit price organization (i.e., located directly next to the product itself on the product shelf). Russo found that the list organization significantly influenced participants purchase decisions by increasing the number of products purchased with lower unit prices. The list organization reduced the cognitive effort needed to acquire and combine unit price information, thus influencing which products people chose to purchase.

Organization by alternatives and organization by attributes are two different yet commonly used types of information organization. Organization by alternatives is typically manipulated in multi-attribute choice tasks by displaying all of the attribute values associated with a single alternative and separate displays for each alternative. In contrast, organization by attributes involves displaying the values of an attribute for multiple alternatives, with separate displays for each attribute.

Organization can determine how people acquire decision information, which can shape how alternatives are evaluated. In general, organization by alternatives leads to alternative-based processing, in which people first acquire information about all of the attributes associated with a single alternative before acquiring information about subsequent alternatives (Payne et al., 1994). Organization by alternatives tends to highlight differences between the alternative's attributes rather than between the alternatives themselves, which can enhance perception of each alternative's stronger attributes.

Using tables, Chang and Liu (2008a) demonstrated how organization influences how people process information which can lead to systematic differences in terms of choice among alternatives. Chang and Liu compared alternative and attribute organizations in a preference-based choice task in which participants chose from a set of consumer products to purchase (e.g., digital cameras) which were described by two attributes (e.g., reliability and picture quality).

The results from Chang and Liu (2008a) showed that participants chose the compromise option (i.e., alternative with the same value for both attributes) more often when information was organized by alternatives than when organized by attributes. Furthermore, participants chose the middle alternative (i.e., the alternative presented second in the list) more often when information was organized by attributes. In effect, organization influenced which option appeared to be the true compromise. Chang and Liu explained this organization effect on choice by extending previous research on compromise effects (e.g., Chernev, 2004; Simonson, 1989). When organized by alternative, the true compromise option (i.e., same value for both attributes) was more salient than the middle

option (i.e., the alternative presented second in the list). In contrast, when organized by attribute, the middle option appeared to be the compromise.

In summary, organization constitutes a fundamental characteristic of an information display that can be manipulated to influence people to engage in alternative or attribute-based processing which can influence choice. Consistent with the concreteness principle (Slovic, 1972), people tend to process decision information in a manner that is consistent with the organization. Organization by alternative leads to alternative-based processing in which people are more likely to process information across attributes to form a holistic evaluation of the alternative. In contrast, organization by attributes leads to attribute-based processing or within-attribute processing, which can bias evaluation of alternatives based on a single attribute as opposed to consideration of all attributes (i.e., examining fewer cues and integrating less information).

8.3 Sequence

Information sequence refers to the ordering of decision and attribute information within the display. Decision information can be sequenced in a number of ways including alphabetical order of the alternative's names or by descending attribute values. For example, cell phones could be sequenced in a table by descending values in storage capacity. Manipulating the sequence of decision information can directly impact the order in which people acquire and process the information. As a result, sequence effects can significantly change how people weigh the information. For instance, sorting cell phones

by descending storage quality can subtly change how people weigh the storage capacity attribute in their evaluations of the cell phones.

Carlson, Meloy, and Russo (2006) demonstrated how the sequence in which attributes were presented could be manipulated to influence choice in favor of one alternative. In study 1 of Carlson et al., participants chose between two consumer products (e.g., backpacks) that were described by six attributes appearing in a table format. Two attributes were diagnostic (i.e., one favored alternative A, one favored alternative B), and the remaining attributes were neutral (i.e., did not favor an alternative). The sequence of the attributes was manipulated such that the first and fourth attributes in the order were diagnostic and the other attributes were neutral. After reviewing all attributes, participants chose one alternative they preferred. Analysis of the choice frequency data showed that the sequence of the diagnostic attributes significantly influenced participants' choices between two alternatives. Specifically, whichever alternative the first diagnostic attribute favored (i.e., the first attribute participants saw) was chosen about 70% of the time. In other words, when the first attribute that participants reviewed significantly favored alternative A, participants were more likely to choose alternative A.

Information sequence can influence choice by changing how people weight decision information that is processed early in the sequence, which can frame how subsequent information is evaluated. Carlson and colleagues (2006) proposed that by placing an attribute that favored one alternative early in the sequence instilled a leader-driven primacy effect. Sequence manipulations can determine the order in which attributes are acquired and processed, which can influence evaluations that occur downstream. As a

result, participants' evaluations of subsequent attributes were biased to favor the alternative that was supported by the first attribute in the sequence. By carefully sequencing attributes and positioning an attribute that favors one alternative early in the sequence, designers can increase the likelihood that people will choose that alternative.

Russo, Carlson, and Meloy (2006) further examined how attribute sequence can even influence people to choose an inferior alternative in terms of its attribute utility value. Using a choice task paradigm and sequence manipulation similar to that of Carlson et al. (2006), Russo and colleagues traced the decision process and were able to find converging evidence that the leader-driven primacy effect does indeed influence choice by changing how attributes are weighted which shapes how participants evaluate subsequent decision information. Assuming that preferences are constructed in the moment (Bettman et al., 1998; Payne et al., 2000), subsequent attribute information is evaluated in a manner that is biased towards supporting the leading alternative. The sequencing effects also suggest the reduction of cognitive effort by simplifying the weighting of attribute information during information combination (Shah & Oppenheimer, 2008). That is, reducing the cognitive operations involved in making trade-offs among the alternatives, strategically manipulating sequence can lead to the observed systematic biases in choice.

In addition to attribute sequence effects, manipulating the order in which alternatives are presented can influence choice (Cai & Xu, 2008). Participants browsed a simulated retail website that included a tabular display containing nine digital cameras, each of which were described by price as well as six attributes pertaining to quality (e.g., number of megapixels). Sequence was manipulated by sorting alternatives based on an

overall quality attribute either in ascending, descending, or random order. Results showed that sequencing alternatives by descending quality led participants to select cameras that were of higher quality than when the alternatives were ordered by ascending quality or randomized. Moreover, sequencing alternatives by descending quality led participants to rate quality as a more important factor than price in their decision. Cai and Xu (2008) explain the effect of descending quality on choice by using the principle of concreteness (Slovic, 1972) and loss aversion (Kahneman & Tversky, 1979). In accordance with the concreteness principle, participants used the quality attribute information in the manner in which it was displayed; therefore, descending quality led people to value quality as more important when evaluating alternatives and including them in their consideration set. Furthermore, with a descending order of quality, higher quality alternatives are established as a reference point at which subsequent comparisons of alternatives are made. Therefore, people choose superior quality alternatives because quality is promoted as a distinguishing factor and because subsequent alternatives are evaluated as a loss in quality (due to loss aversion).

In summary, sequence influences choice by determining the order in which information is likely to be processed, which can impact how subsequent attributes are weighted when evaluating alternatives. As a fundamental characteristic of the information display, sequence can be manipulated in any multi-attribute choice decision context, regardless of format or organization.

8.4 Salience

The salience category of display factors refers to perceptual manipulations (as opposed to cognitive or internal salience) that change the perceptual salience of specific alternatives or attributes. In general, salience manipulations influence decision processes by impacting how people allocate their selective visual attention towards specific pieces of decision information. According to Wolfe's (2007) Guided Search Model, selective visual attention is influenced by bottom-up activation (i.e., local differences in perceptual values) and top-down activation (i.e., similarity of visual stimuli to an internal representation). Manipulating the salience of decision information offers a means of influencing such bottom-up activation levels. Salience effects on selective attention can influence whether people will attend to a specific decision alternative or attribute and the duration for which they do so. The consequence of which ultimately determine whether people will choose one alternative over another (e.g., Towal, Mormann, & Koch, 2013).

In the human factors literature, researchers have been interested in how manipulating salience can influence selective visual attention and aid decision-making performance (e.g., Bennett, Naggy, & Flach, 2006; Nikolic, Orr, & Sarter, 2004; Wickens, McCarley, & Steelman-Allen, 2009). For example, Nikolic and colleagues showed how manipulating the visual salience of alert information by altering location, color, and movement can increase the likelihood that operators will attend to and process the alert, which can lead to improve system performance. Research from the HAI literature suggests that manipulating salience can play an important role in improving HAI outcomes. In an effort to increase automation transparency and reduce automation biases, Rovira et al., (2014) used a color-coded highlighting scheme to convey changes in the reliability of the

automation over time. They found that manipulating salience of the automation's reliability increased appropriate usage of the automation by participants. These studies illustrate how human factors researchers have applied salience effects to address visual display design problems and improve performance. However, these results are limited when considering how to engineer a decision environment for a multi-attribute choice decision-making context.

In the consumer decision-making literature, researchers have investigated how salience manipulations can lead to a visual bias effect on choice. The visual bias effect refers to the finding that the amount of relative visual attention allocated towards a given alternative can serve as an important predictor of people's choice among alternatives (Armel, Beaumel, & Rangel, 2008; Krajbich & Rangel, 2011; Pieters & Warlop, 1997; Towal et al., 2013). For example, Armel et al. (2008) manipulated the presentation duration of alternatives and found that people generally preferred the alternative that was presented for a longer duration. This suggests that increasing time spent examining an option will influence the individual's preference for that option. In an eye tracking study by Pieters and Warlop (1997), eye movement measures were used to successfully predict participants' choices. Participants fixated chosen alternatives more often than non-chosen alternatives and participants skipped fixating fewer visual elements associated with the chosen alternative. Additional eye tracking studies provide further support for the visual bias effect on choice and extend it to include fixation duration as a predictive measure of choice (Krajbich & Rangel, 2011; Towal et al., 2013). Thus, the relative visual attention allocated

towards a given alternative constitutes an important factor in the decision-making process, which can determine whether an individual will choose a particular alternative.

In an effort to leverage the visual bias effect on choice, researchers have explored how the relative brightness of alternatives can be manipulated to increase salience. Milosavljevic, Navalpakkam, Koch, and Rangel (2012) manipulated the relative brightness of an alternative displayed to participants by decreasing the brightness of other alternatives in the display by 65%. Results showed the salience manipulation significantly increased the likelihood that people would choose the salient alternative. These results can be explained in terms of a visual bias effect on the allocation of selective visual attention in favor of the salient alternative.

In addition to relative brightness, researchers have examined other methods of increasing perceptual salience of decision information including highlighting text and increasing font size. Specifically, Jiang and Punj (2010) highlighted attribute information (i.e., restaurant atmosphere ratings and price) in green and increased the font size to enhance salience. The results showed that in contrast to the non-salient attribute, participants examined more alternatives that were high in value for the salient attribute and increased the likelihood that participants would choose an alternative that was superior in terms of the salient attribute. Therefore, enhancing attribute salience appears to lead people to acquire information and process alternatives based on the salient attribute. This interpretation is similar to that of the previously discussed visual bias effect in that highlighting an attribute can increase the relative visual attention allocated to those attributes. Consequently, highlighting attributes biases selective visual attention, framing

how the attributes are processed and evaluated, which influences choice in favor of the alternatives that are high on the salient attribute.

To summarize, salience is a fundamental characteristic of an information display that can be manipulated to influence the allocation of selective visual attention and change how people evaluate alternatives. The visual bias effect shows that the way people allocate their visual attention can shape how decision alternatives are evaluated. By manipulating the information salience in the design of a decision environment, designers can tap into the visual bias effect in predictable ways by increasing the likelihood that people allocate their visual attention to the salient information. Although there are numerous methods available for designers to increase salience, the efficacy of different salience factors is unclear due to limited research investigating salience effects on choice.

8.5 Interactions of Display Factors

Designing a visual display of decision information involves instantiating format, organization, sequence, and salience. Research investigating how display factors combine to impact decision-making is limited but there is some evidence to indicate that certain combinations of display factors can minimize their influence on choice. Therefore, it is crucial to consider the confluence of display factors on decision making processes and choice.

Considering the interaction of sequence and salience, there is some evidence to suggest that manipulating the sequence of alternatives can attenuate salience effects on choice (Jiang & Punj, 2010). Specifically, salience was manipulated by highlighting and

increasing the font size associated with one of two attributes; sequence was manipulated by ordering alternatives either by name (in alphabetical order) or by descending values for the non-salient attribute. The results showed that when alternatives were sequenced by name, choice share for alternatives that were high on the salient attribute increased. However, when alternatives were sequenced by the non-salient attribute, there was no difference in choice share between alternatives that were high on the salient attribute and those that were high on the non-salient attribute. Thus, the combination of the salience and sequence manipulations cancelled out the potentially influential effects of each manipulation.

Organization has been shown to moderate the influence of sequence (Carlson et al., 2006). In study 3 of Carlson et al., sequence was manipulated in a multi-attribute binary choice task by ordering attributes such that the first attribute significantly favored one alternative, the fourth attribute favoring the other alternative, and the remaining attributes constituting neutral attributes. Information organization was manipulated by organizing the information either by attribute (i.e., simultaneous display of attributes for both alternatives) or by alternative (i.e., separate displays; one containing attributes for alternative A and one for alternative B). Analysis of participants' choices revealed a significant difference between the attribute and alternative organization conditions. For the attribute organization, 86% of participants chose the alternative that was favored by the first attribute. However, for the alternative organization condition, only 58% of participants chose the alternative which the first attribute favored. This result indicates that

manipulating information organization reduced the leader-driven primacy effect that was attributed to the sequence manipulation.

These studies illustrate the importance of considering potential interactions of display factors. Furthermore, due to the limited available research, it remains unclear how display factors can be combined to enhance their relative influence on choice. For example, Jiang and Punj (2010) showed that sequence can minimize salience effects when each factor is in conflict, but it is unclear if both factors encouraged processing of the same attributes whether that would strengthen their influence in favor of alternatives that were high on that sequenced and salient attribute. By considering the potential interactions, designers can better predict how a given display design will impact decision making. Furthermore, understanding how a given configuration of display factors impacts decision making is critical for applying display design effects to improve real world decision making like that involved in human-automation interaction. After all, designing a display for an automated decision support system involves the interaction of numerous display factors.

APPENDIX B. TASK INSTRUCTIONS FOR EXPERIMENT 1

Introduction

Hello and thank you for agreeing to participate in this experiment. **Please carefully read the following instructions before you begin the experiment**.

In this experiment, you will take on the role of a package delivery service agent. Your task will be to navigate to various destinations to deliver packages. To deliver each package, you will choose one of four different routes for each of your deliveries. After choosing a route, you will navigate to your destination by monitoring your vehicle's position along the route. During transit, you will **not** control the vehicle, but you need to make sure that each package is delivered. After delivering the package, you will choose from a new set of routes in order to deliver the next package to the next destination.

Each delivery will involve the following procedure:

1. Review 4 different routes, each displayed on a separate page
2. Choose your preferred route from the 4 routes
3. Monitor your vehicle's position as it navigates to the destination
4. Await confirmation that the package has been delivered
5. Begin next delivery

Press NEXT to learn more about how the routes are presented.

Routes

Below is an example of a single route display with the four route attributes. Note each part of the route display as each delivery will require that you review four routes presented in this same format.

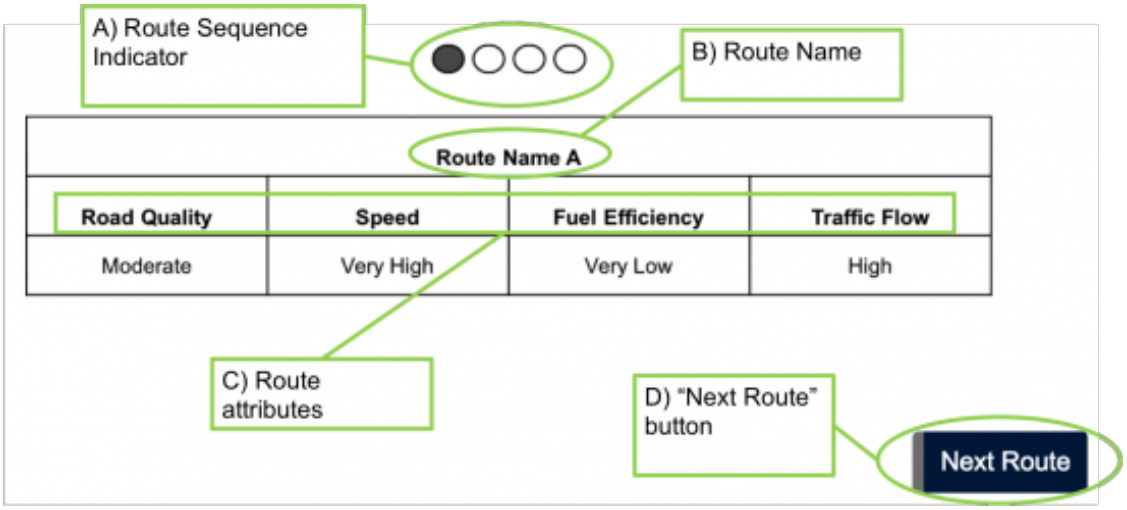

A) Route sequence indicator

- Only one route will be displayed per page and this indicator depicts which route in the set of 4 is being displayed.

B) Route name

- Each route will have a unique name and the route name identifies each route.

C) Route attributes

- Depicts the name of each attribute with the correspond value directly below that attribute. Route attributes can range in value from:
    - Very Low - Low
    - Low - Moderate
    - Moderate - High
    - High - Very High

D) Next Route button

- Advances to the next route in the sequence of routes.

<u>Route Attributes</u>

Each route will be described by 4 attributes and each attribute will range in value from "Very Low" to "Very High". All routes are described by each of these four attributes. You should consider these attributes as you evaluate and choose a route. It is important that you only consider the attributes as they are defined below. **Please carefully read each attribute definition before beginning the experiment**.

Fuel Efficiency

- The level of fuel efficiency for your vehicle along the route.
- "Very Low" fuel efficiency indicates that your vehicle will have poor fuel mileage along that route because your vehicle must work harder to travel to the destination.
- "Very High" fuel efficiency indicates that your vehicle runs at optimal efficiency along this route, thereby using less fuel to travel along this route.

Road Quality

- Road quality refers to the actual road or street conditions and how the road conditions will impact your vehicle's operating condition, your comfort from within the cabin of the vehicle, as well as the cargo you're delivering.
- "Very Low" road quality indicates that your vehicle is more likely to be damaged. You and your cargo are also more likely to be jostled or moved about within the vehicle as you navigate to your destination.
- "Very High" road quality indicates that you are unlikely to encounter poor road conditions that could damage your vehicle and jostle or move you and the cargo within the vehicle.

Speed

- Speed is determined by the average speed of your vehicle along the route, which is due to the speed limit on the roadways, in addition to other factors.
- "Very Low" speed indicates that your vehicle will travel slowly along the route.
- "Very High" speed indicates that your vehicle will travel quickly along the route.

Traffic Flow

- Traffic flow refers to the efficiency with which traffic is moving along the route.
- "Very Low" traffic flow indicates that traffic is not flowing smoothly due to higher traffic density and congestion, in addition to other factors.
- "Very High" traffic flow indicates that traffic is flowing smoothly because there is low traffic density and less congestion along this route.

Please note that Speed and Traffic Flow do **NOT** necessarily determine how quickly you will arrive at your destination because the estimated time of arrival and the duration of travel along each route are unknown by the system.

Remember, there are no right or wrong answers. We want you to choose the route that you prefer, but you should only consider these 4 attributes when you're

evaluating and choosing a route.

As a delivery service agent, it is important that you deliver the packages in a timely fashion. However, it is equally important that you choose the routes that allow your vehicle to operate efficiently, and that your vehicle and the cargo are not damaged.

You will notice that some route displays may have some attribute information highlighted in yellow. The purpose for highlighting information is to determine whether this might help people make their decisions. The highlighted information does not indicate that it is any more or less important. Please use this highlighted information however you would like.

Examples

Below are some pictures that depict the sequence of events for each package you will deliver.

1. You will first review the 4 routes on separate pages.
2. After you review a given route, you can press the "Next Route" button to begin viewing the next route's attributes.
3. After reviewing all 4 routes, you will be able to choose your preferred route among all 4 routes.

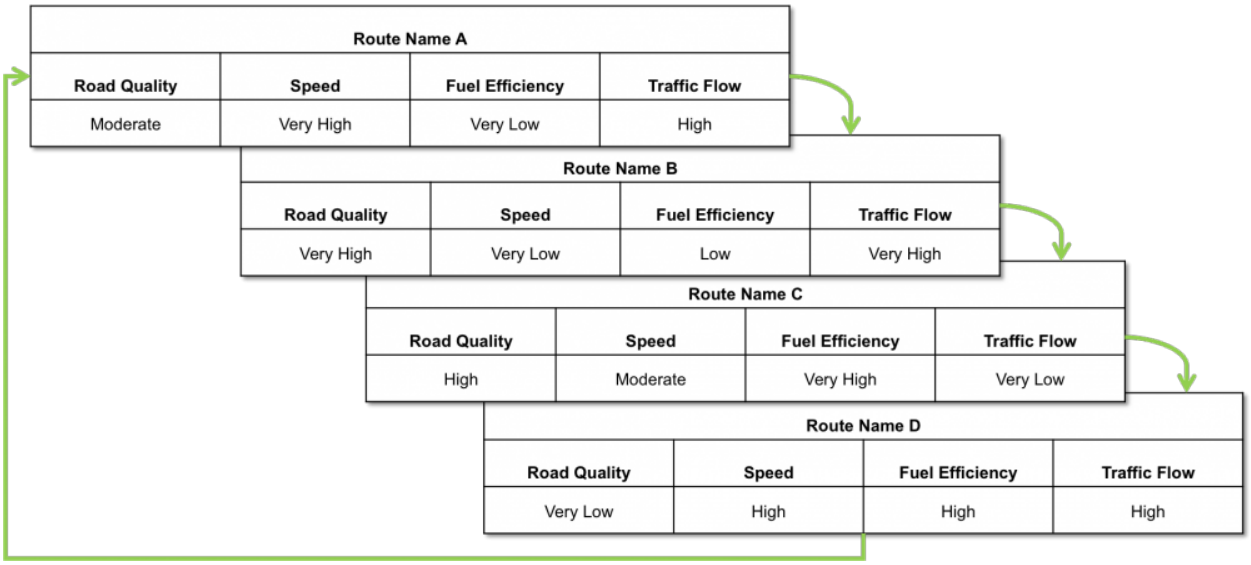

During the first time through each set of 4 routes, you will **not** be able to choose a route. Just press the "Next Route" button to advance to the next route.

After you have reviewed the 4th route, press the "Next" button.

At this point, you will be able to choose 1 of the 4 routes by clicking on the radio button next to your preferred route:

Choose your preferred route.

Route Name A			
Road Quality	**Speed**	**Fuel Efficiency**	**Traffic Flow**
Moderate	Very High	Very Low	High

After choosing a route, you will view the GPS route system display. Your vehicle's location will be displayed as a blue circle and your destination will be marked as a red pentagon. Your task is to monitor your vehicle's position along the route to ensure that the package is delivered.

APPENDIX C. TASK INSTRUCTIONS FOR EXPERIMENT 3

Automated System and Route Recommendations

For each delivery, an automated system will recommend 1 of the 4 routes for you to choose. The system's recommended route is based on an evaluation of each route's projected attribute values as well as other information such as weather forecast.

Through extensive testing, we have determined that ***the automation's recommendations are only 70% reliable***. This means that 70% of the time, the automation recommended a route which was optimal or equivalent to other high performing routes. On the other hand, 30% of the time, the recommended route was sub-optimal and performed poorly in comparison to one or more of the other routes that were available. The system will only recommend a route; it will not automatically choose a route for you. For each delivery, the first of the 4 routes will be the systems recommended route, followed by 3 alternative routes.

You should consider the system's recommended route, but you should also evaluate all the other routes and choose the one you believe will lead to the best performance.

Again, your goal as a delivery service agent is to choose a route that achieves the following:

- Deliver cargo/packages in a timely fashion
- Ensure your vehicle will operate efficiently along the route
- Your vehicle and the cargo/packages are not damaged during transport due to route conditions

Finally, you will notice that some of the attributes are highlighted in yellow. These highlighted attributes are chosen by the automated system and convey why the system is recommending that route. You should use this highlighted information however you prefer.

APPENDIX D. TASK INSTRUCTIONS FOR EXPERIMENT 4

Automated System and Route Recommendations

For each delivery, an automated system will recommend 1 of the 4 routes for you to choose. The system's recommended route is based on an evaluation of each route's projected attribute values as well as other information such as weather forecast.

Through extensive testing, we have determined that ***the automation's recommendations are 90% reliable***. This means that 90% of the time, the automation recommended a route which was optimal or equivalent to other high performing routes. On the other hand, 10% of the time, the recommended route was sub-optimal and performed poorly in comparison to one or more of the other routes that were available. The system will only recommend a route; it will not automatically choose a route for you. For each delivery, the first of the 4 routes will be the systems recommended route, followed by 3 alternative routes.

You should consider the system's recommended route, but you should also evaluate all the other routes and choose the one you believe will lead to the best performance.

You will notice that sometimes an attribute will be highlighted in yellow. These highlighted attributes are chosen by the automated system and are intended to convey why the system is recommending that route. However, this does not mean that the system's recommendation is any more or less valid. Please use this highlighted information however you would like to help you choose your route.

Again, your goal as a delivery service agent is to choose a route that achieves the following:

- Deliver cargo/packages in a timely fashion
- Ensure your vehicle will operate efficiently along the route
- Your vehicle and the cargo/packages are not damaged during transport due to route conditions

APPENDIX E. ROUTE ATTRIBUTE UTILITY VALUES FOR EXPERIMENTS 1 AND 2

Table 4 – Route Attribute Utility Values for Experiments 1 and 2

Trial	Route Set	Route	Target Route	Seq.	Sal.	Target Attrib.	Road Qual.	Speed	Fuel Effic.	Traffic Flow	Utility
1	4	A	D	Yes	Yes	Speed	3	4	1	5	13
1	4	B	D	Yes	Yes	Speed	5	1	3	4	13
1	4	C	D	Yes	Yes	Speed	4	3	5	1	13
1	4	D	D	Yes	Yes	Speed	1	5	4	3	13
2	12	A	D	Yes	No	Fuel Efficiency	2	5	2	4	13
2	12	B	D	Yes	No	Fuel Efficiency	4	2	2	5	13
2	12	C	D	Yes	No	Fuel Efficiency	5	2	4	2	13
2	12	D	D	Yes	No	Fuel Efficiency	2	4	5	2	13
3	23	A	A	No	Yes	Traffic Flow	1	4	1	5	11
3	23	B	A	No	Yes	Traffic Flow	5	1	1	4	11
3	23	C	A	No	Yes	Traffic Flow	4	1	5	1	11
3	23	D	A	No	Yes	Traffic Flow	1	5	4	1	11
4	26	A	C	No	No	Road Quality	3	5	4	4	16
4	26	B	C	No	No	Road Quality	4	4	3	5	16
4	26	C	C	No	No	Road Quality	5	3	4	4	16
4	26	D	C	No	No	Road Quality	4	4	5	3	16
5	4	A	D	Yes	No	Speed	3	4	1	5	13
5	4	B	D	Yes	No	Speed	5	1	3	4	13
5	4	C	D	Yes	No	Speed	4	3	5	1	13
5	4	D	D	Yes	No	Speed	1	5	4	3	13
6	12	A	D	No	Yes	Fuel Efficiency	2	5	2	4	13
6	12	B	D	No	Yes	Fuel Efficiency	4	2	2	5	13

6	12	C	D	No	Yes	Fuel Efficiency	5	2	4	2	13
6	12	D	D	No	Yes	Fuel Efficiency	2	4	5	2	13
7	23	A	A	No	No	Traffic Flow	1	4	1	5	11
7	23	B	A	No	No	Traffic Flow	5	1	1	4	11
7	23	C	A	No	No	Traffic Flow	4	1	5	1	11
7	23	D	A	No	No	Traffic Flow	1	5	4	1	11
8	26	A	C	Yes	Yes	Road Quality	3	5	4	4	16
8	26	B	C	Yes	Yes	Road Quality	4	4	3	5	16
8	26	C	C	Yes	Yes	Road Quality	5	3	4	4	16
8	26	D	C	Yes	Yes	Road Quality	4	4	5	3	16
9	4	A	D	No	Yes	Speed	3	4	1	5	13
9	4	B	D	No	Yes	Speed	5	1	3	4	13
9	4	C	D	No	Yes	Speed	4	3	5	1	13
9	4	D	D	No	Yes	Speed	1	5	4	3	13
10	12	A	D	No	No	Fuel Efficiency	2	5	2	4	13
10	12	B	D	No	No	Fuel Efficiency	4	2	2	5	13
10	12	C	D	No	No	Fuel Efficiency	5	2	4	2	13
10	12	D	D	No	No	Fuel Efficiency	2	4	5	2	13
11	23	A	A	Yes	Yes	Traffic Flow	1	4	1	5	11
11	23	B	A	Yes	Yes	Traffic Flow	5	1	1	4	11
11	23	C	A	Yes	Yes	Traffic Flow	4	1	5	1	11
11	23	D	A	Yes	Yes	Traffic Flow	1	5	4	1	11
12	26	A	C	Yes	No	Road Quality	3	5	4	4	16
12	26	B	C	Yes	No	Road Quality	4	4	3	5	16
12	26	C	C	Yes	No	Road Quality	5	3	4	4	16
12	26	D	C	Yes	No	Road Quality	4	4	5	3	16
13	4	A	D	No	No	Speed	3	4	1	5	13
13	4	B	D	No	No	Speed	5	1	3	4	13

13	4	C	D	No	No	Speed	4	3	5	1	13
13	4	D	D	No	No	Speed	1	5	4	3	13
14	12	A	D	Yes	Yes	Fuel Efficiency	2	5	2	4	13
14	12	B	D	Yes	Yes	Fuel Efficiency	4	2	2	5	13
14	12	C	D	Yes	Yes	Fuel Efficiency	5	2	4	2	13
14	12	D	D	Yes	Yes	Fuel Efficiency	2	4	5	2	13
15	23	A	A	Yes	No	Traffic Flow	1	4	1	5	11
15	23	B	A	Yes	No	Traffic Flow	5	1	1	4	11
15	23	C	A	Yes	No	Traffic Flow	4	1	5	1	11
15	23	D	A	Yes	No	Traffic Flow	1	5	4	1	11
16	26	A	C	No	Yes	Road Quality	3	5	4	4	16
16	26	B	C	No	Yes	Road Quality	4	4	3	5	16
16	26	C	C	No	Yes	Road Quality	5	3	4	4	16
16	26	D	C	No	Yes	Road Quality	4	4	5	3	16

APPENDIX F. ROUTE ATTRIBUTE UTILITY VALUES FOR EXPERIMENTS 3 AND 4

Table 5 – Route Attribute Utility Values for Experiments 3 and 4

Trial #	Trial Type	Route Set	Route	Rec. Route	Targeted Attribute	Road Qual.	Speed	Fuel Effic.	Traffic Flow	Utility
1	Equal	1	A	A	Road Quality	5	1	4	2	12
1	Equal	1	B	A	Road Quality	2	4	1	5	12
1	Equal	1	C	A	Road Quality	4	5	2	1	12
1	Equal	1	D	A	Road Quality	1	2	5	4	12
2	Equal	2	A	B	Speed	1	3	4	5	13
2	Equal	2	B	B	Speed	3	5	1	4	13
2	Equal	2	C	B	Speed	5	4	3	1	13
2	Equal	2	D	B	Speed	4	1	5	3	13
3	Equal	3	A	C	Fuel Efficiency	3	5	2	1	11
3	Equal	3	B	C	Fuel Efficiency	2	3	1	5	11
3	Equal	3	C	C	Fuel Efficiency	1	2	5	3	11
3	Equal	3	D	C	Fuel Efficiency	5	1	3	2	11
4	Equal	4	A	D	Traffic Flow	1	2	4	3	10
4	Equal	4	B	D	Traffic Flow	4	1	3	2	10
4	Equal	4	C	D	Traffic Flow	3	4	2	1	10
4	Equal	4	D	D	Traffic Flow	2	3	1	4	10
5	Superior	5	A	A	Road Quality	5	3	3	2	13
5	Superior	5	B	A	Road Quality	3	2	3	4	12
5	Superior	5	C	A	Road Quality	2	3	4	3	12
5	Superior	5	D	A	Road Quality	3	4	1	4	12
6	Superior	6	A	B	Speed	4	2	3	3	12

6	Superior	6	B	B	Speed	3	5	2	3	13
6	Superior	6	C	B	Speed	3	2	4	3	12
6	Superior	6	D	B	Speed	1	4	3	4	12
7	Superior	5	A	A	Road Quality	5	3	3	2	13
7	Superior	5	B	A	Road Quality	3	2	3	4	12
7	Superior	5	C	A	Road Quality	2	3	4	3	12
7	Superior	5	D	A	Road Quality	3	4	1	4	12
8	Superior	6	A	B	Speed	4	2	3	3	12
8	Superior	6	B	B	Speed	3	5	2	3	13
8	Superior	6	C	B	Speed	3	2	4	3	12
8	Superior	6	D	B	Speed	1	4	3	4	12
9	Equal	11	A	C	Fuel Efficiency	3	5	2	1	11
9	Equal	11	B	C	Fuel Efficiency	2	3	1	5	11
9	Equal	11	C	C	Fuel Efficiency	1	2	5	3	11
9	Equal	11	D	C	Fuel Efficiency	5	1	3	2	11
10	Equal	12	A	D	Traffic Flow	1	2	4	3	10
10	Equal	12	B	D	Traffic Flow	4	1	3	2	10
10	Equal	12	C	D	Traffic Flow	3	4	2	1	10
10	Equal	12	D	D	Traffic Flow	2	3	1	4	10
11	Equal	11	A	C	Fuel Efficiency	3	5	2	1	11
11	Equal	11	B	C	Fuel Efficiency	2	3	1	5	11
11	Equal	11	C	C	Fuel Efficiency	1	2	5	3	11
11	Equal	11	D	C	Fuel Efficiency	5	1	3	2	11
12	Equal	12	A	D	Traffic Flow	1	2	4	3	10
12	Equal	12	B	D	Traffic Flow	4	1	3	2	10
12	Equal	12	C	D	Traffic Flow	3	4	2	1	10

12	Equal	12	D	D	Traffic Flow	2	3	1	4	10
13	Inferior	13	A	A	Road Quality	5	3	2	1	11
13	Inferior	13	B	A	Road Quality	1	3	4	4	12
13	Inferior	13	C	A	Road Quality	3	3	4	2	12
13	Inferior	13	D	A	Road Quality	2	4	2	4	12
14	Inferior	14	A	B	Speed	4	3	4	1	12
14	Inferior	14	B	B	Speed	2	5	1	3	11
14	Inferior	14	C	B	Speed	4	3	3	2	12
14	Inferior	14	D	B	Speed	2	2	4	4	12
15	Inferior	13	A	A	Road Quality	5	3	2	1	11
15	Inferior	13	B	A	Road Quality	1	3	4	4	12
15	Inferior	13	C	A	Road Quality	3	3	4	2	12
15	Inferior	13	D	A	Road Quality	2	4	2	4	12
16	Inferior	14	A	B	Speed	4	3	4	1	12
16	Inferior	14	B	B	Speed	2	5	1	3	11
16	Inferior	14	C	B	Speed	4	3	3	2	12
16	Inferior	14	D	B	Speed	2	2	4	4	12

REFERENCES

Arunachalam, V., Pei, B. K., & Steinbart, P. J. (2002). Impression management with graphs: Effects on choices. *Journal of Information Systems*, *16*(2), 183-202.

Armel, K. C., Beaumel, A., & Rangel, A. (2008). Biasing simple choices by manipulating relative visual attention. *Judgment and Decision Making*, *3*, 396–403.

Bertin, J. (1983). Semiology of graphics: Diagrams networks maps (W. Berg, Trans.). Madison: University of Wisconsin Press.

Bettman, J. R., Johnson, E. J., & Payne, J. W. (1991). Consumer decision making. *Handbook of Consumer Behavior, 6*, 50-84.

Bettman, J. R., & Kakkar, P. (1977). Effects of information presentation format on consumer information acquisition strategies. *Journal of Consumer Research*, *3*(4), 233.

Bettman, J. R., Luce, M. F., & Payne, J. W. (1998). Constructive consumer choice processes. *Journal of consumer research*, *25*(3), 187-217.

Cai, S., & Xu, Y. (2008). Designing Product Lists for E-commerce: The Effects of Sorting on Consumer Decision Making. *International Journal of Human-Computer Interaction*, *24*(7), 700–721.

Carlson, K. A., Meloy, M. G., & Russo, J. E. (2006). Leader-driven primacy: Using attribute order to affect consumer choice. Journal of Consumer Research, 32(4), 513-518.

Chang, C. C., & Liu, H. H. (2008a). Information format-option characteristics compatibility and the compromise effect. *Psychology & Marketing*, *25*(9), 881-900.

Chang, C. C., & Liu, H. H. (2008b). Which is the compromise option? Information format and task format as determinants. *Journal of Behavioral Decision Making*, *21*(1), 59-75.

Chen, J. Y., Procci, K., Boyce, M., Wright, J., Garcia, A., & Barnes, M. (2014). *Situation awareness-based agent transparency* (No. ARL-TR-6905). Army Research Lab Aberdeen Proving Ground, MD Human Research and Engineering Directorate.

Chernev, A. (2004). Extremeness aversion and attribute–balance effects in choice. Journal of Consumer Research, 31, 249–263.

de Visser, E. J., Pak, R., & Shaw, T. H. (2018). From 'automation' to 'autonomy': the importance of trust repair in human–machine interaction. *Ergonomics*, 1-19

Dhar, R., & Gorlin, M. (2013). A dual-system framework to understand preference construction processes in choice. *Journal of Consumer Psychology*, *23*(4), 528-542.

Dhar, R., Nowlis, S. M., & Sherman, S. J. (2000). Trying hard or hardly trying: An analysis of context effects in choice. *Journal of Consumer Psychology*, *9*(4), 189-200.

Dilla, W.N., & Steinbart, P.J. (2005) Using information display characteristics to provide decision guidance in a choice task under conditions of strict uncertainty. *Journal of information systems, 19,* 29-55.

Dorneich, M. C., Dudley, R., Letsu-Dake, E., Rogers, W., Whitlow, S. D., Dillard, M. C., & Nelson, E. (2017). Interaction of Automation Visibility and Information Quality in Flight Deck Information Automation. *IEEE Transactions on Human-Machine Systems*, *47*(6), 915-926.

Durso, F. T., & Alexander, A. L. (2010). Managing workload, performance, and situation awareness in aviation systems. In *Human Factors in Aviation (Second Edition)* (pp. 217-247).

Endsley, M. R. Toward A Theory of Situation Awareness in Dynamic Systems. Human Factors 1995, 37 (1), 32–64.

Felfernig, A., Friedrich, G., Gula, B., Hitz, M., Kruggel, T., Leitner, G., ... & Vitouch, O. (2007). Persuasive recommendation: serial position effects in knowledge-based recommender systems. In Persuasive Technology (pp. 283-294). Springer Berlin Heidelberg.

Fennema, M. G., & Kleinmuntz, D. N. (1995). Anticipations of effort and accuracy in multiattribute choice. *Organizational behavior and human decision processes*, *63*(1), 21-32.

Few, S. (2004). *Show me the numbers: Designing tables and graphs to enlighten* (Vol. 1, No. 1). Oakland, CA: Analytics Press.

Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart.* Oxford, England: Oxford University Press.

Goddard, K., Roudsari, A., & Wyatt, J. C. (2011). Automation bias: a systematic review of frequency, effect mediators, and mitigators. *Journal of the American Medical Informatics Association*, *19*(1), 121-127.

Goddard, K., Roudsari, A., & Wyatt, J. C. (2014). Automation bias: empirical results assessing influencing factors. *International journal of medical informatics*, *83*(5), 368-375.

Hancock, P. A. (2017). Imposing limits on autonomous systems. *Ergonomics*, *60*(2), 284-291.

Hegarty, M. (2011). The cognitive science of visual-spatial displays: Implications for design. *Topics in cognitive science*, *3*(3), 446-474.

Helldin, T. (2014). *Transparency for Future Semi-Automated Systems: Effects of transparency on operator performance, workload and trust* (Doctoral dissertation, Örebro Universitet).

Hoff, K. A., & Bashir, M. (2015). Trust in automation: Integrating empirical evidence on factors that influence trust. *Human Factors*, *57*(3), 407-434.

Janiszewski, C. (1998). The influence of display characteristics on visual exploratory search behavior. *Journal of Consumer Research*, *25*(3), 290–301.

Janiszewski, C., Kuo, A., & Tavassoli, N. T. (2013). The influence of selective attention and inattention to products on subsequent choice. *Journal of Consumer Research*, *39*(6), 1258-1274.

Jarvenpaa, S. L. (1989). The effect of task demands and graphical format on information processing strategies. *Management Science*, *35*, 285–303.

Jarvenpaa, S. L. (1990). Graphic displays in decision making - the visual salience effect. *Journal of Behavioral Decision Making*, *3*(May), 247–262.

Jarvenpaa, S. L., & Dickson, G. W. (1988). Graphics and managerial decision making: Research-based guidelines. *Communications of the ACM*,*31*(6), 764-774.

Johnson, E. J., Shu, S. B., Dellaert, B. G., Fox, C., Goldstein, D. G., Häubl, G., ... & Weber, E. U. (2012). Beyond nudges: Tools of a choice architecture. *Marketing Letters*, *23*(2), 487-504.

Jang, J., Trickett, S. B., Schunn, C. D., & Trafton, J. G. (2012). Unpacking the temporal advantage of distributing complex visual displays. *International Journal of Human-Computer Studies*, *70*(11), 812-827.

Jiang, Y., & Punj, G. N. (2010). The effects of attribute concreteness and prominence on selective processing, choice, and search experience. *Journal of the Academy of Marketing Science*, *38*(4), 471–489.

Johnson, E. J., Shu, S. B., Dellaert, B. G., Fox, C., Goldstein, D. G., Häubl, G., ... & Wansink, B. (2012). Beyond nudges: Tools of a choice architecture. *Marketing Letters*, *23*(2), 487-504.

Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. *Heuristics and biases: The psychology of intuitive judgment*, *49*.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47,* 263–291.

Kelton, A. S., Pennington, R. R., & Tuttle, B. M. (2010). The effects of information presentation format on judgment and decision making: A review of the information systems research. *Journal of Information Systems*, *24*(2), 79–105.

Kim, H. J., & Yoon, S. O. (2015). The effect of category label specificity on consumer choice. Marketing Letters, 1-13.

Kleinmuntz, D. N., & Schkade, D. a. (1993). Information displays and decision processes. *Psychological Science*, *4*(4), 221–227.

Kosslyn, S. M. (2006). *Graph design for the eye and mind*. New York: Oxford University Press.

Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, *108*(33), 13852-13857.

Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive science*, *11*(1), 65-100.

Lee, J. D. (2008). Review of a pivotal Human Factors article: "Humans and automation: use, misuse, disuse, abuse". *Human Factors*, *50*(3), 404-410.

Lee, J. D., & See, K. A. (2004). Trust in automation: designing for appropriate reliance. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *46*(1), 50–80.

Lindberg, E., Gärling, T., & Montgomery, H. (1991). Prediction of preferences for and choices between verbally and numerically described alternatives. *Acta Psychologica*, *76*(2), 165–176.

Lurie, N., & Mason, C. (2007). Visual representation: Implications for decision making. *Journal of Marketing*, *71*, 160–177.

Lyons, J. B. (2013). Being transparent about transparency. In *AAAI Spring Symposium.*

Lyons, J. B., Koltai, K. S., Ho, N. T., Johnson, W. B., Smith, D. E., & Shively, R. J. (2016). Engineering trust in complex automated systems. *ergonomics in design*, *24*(1), 13-17.

Lyons, J. B., Sadler, G. G., Koltai, K., Battiste, H., Ho, N. T., Hoffmann, L. C., ... & Shively, R. (2017). Shaping trust through transparent design: theoretical and experimental guidelines. In *Advances in Human Factors in Robots and Unmanned Systems* (pp. 127-136). Springer, Cham

Mandel, N., & Johnson, E. J. (2002). When web pages influence choice: Effects of visual primes on experts and novices. *Journal of Consumer Research*, 29(2), 235-245.

Mantel, S. P., & Kardes, F. R. (1999). The role of direction of comparison, attribute-based processing, and attitude-based processing in consumer preference. *Journal of Consumer Research*, *25*(4), 335-352.

Martin, J. M., & Norton, M. I. (2009). Shaping online consumer choice by partitioning the web. Psychology & Marketing, 26(10), 908-926.

Mercado, J. E., Rupp, M. A., Chen, J. Y., Barnes, M. J., Barber, D., & Procci, K. (2016). Intelligent agent transparency in human–agent teaming for Multi-UxV management. *Human factors*, *58*(3), 401-415.

Milosavljevic, M., Navalpakkam, V., Koch, C., & Rangel, A. (2012). Relative visual saliency differences induce sizable bias in consumer choice. *Journal of Consumer Psychology*, *22*(1), 67–74.

Mosier, K. L., & Skitka, L. J. (1996). Human decision makers and automated decision aids: Made for each other? In R. Parasuraman & M. Mouloua (Eds.), *Automation and human performance: Theory and application* (pp. 201–220). Mahwah, NJ: Erlbaum.

Mosier, K. L., Skitka, L. J., Heers, S., & Burdick, M. (1998). Automation bias: Decision making and performance in high-tech cockpits. *The International journal of aviation psychology*, *8*(1), 47-63.

Mukherjee, S., & Srinivasan, N. (2013). Attention in preferential choice. Prog. Brain Res, 202, 117-134.

Münscher, R., Vetter, M., & Scheuerle, T. (2016). A review and taxonomy of choice architecture techniques. *Journal of Behavioral Decision Making*, *29*(5), 511-524.

Nikolic, M. I., Orr, J. M., & Sarter, N. B. (2004). Why pilots miss the green box: How display context undermines attention capture. *The International Journal of Aviation Psychology*, *14*(1), 39-52.

O'Donnell, E., & David, J. S. (2000). How information systems influence user decisions: a research framework and literature review. *International Journal of Accounting Information Systems*, *1*(3), 178-203.

Ososky, S., Sanders, T., Jentsch, F., Hancock, P., & Chen, J. Y. (2014, June). Determinants of system transparency and its influence on trust in and reliance on unmanned robotic systems. In *Unmanned Systems Technology XVI* (Vol. 9084, p. 90840E). International Society for Optics and Photonics.

Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. *Human factors*, *52*(3), 381-410.

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on systems, man, and cybernetics-Part A: Systems and Humans*, *30*(3), 286-297.

Parasuraman, R., & Wickens, C. D. (2008). Humans: Still vital after all these years of automation. *Human Factors*, *50*, 511–520.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge University Press.

Payne, J. W., Bettman, J. R., Schkade, D. A., Schwarz, N., & Gregory, R. (1999). Measuring constructed preferences: Towards a building code. In *Elicitation of preferences* (pp. 243-275). Springer, Dordrecht

Pieters, R., & Warlop, L. (1997). Visual attention during brand choice: an eye-fixation analysis. *DTEW Research Report 9724*, 1–41.

Rasmussen, J., & Vicente, K. J. (1989). Coping with human errors through system design: implications for ecological interface design. *International Journal of Man-Machine Studies*, *31*(5), 517-534.

Renshaw, J. A., Finlay, J. E., Tyfa, D., & Ward, R. D. (2004). Understanding visual influence in graph design through temporal and spatial eye movement characteristics. *Interacting with computers*, *16*(3), 557-578.

Rovira, E., Cross, A., Leitch, E., & Bonaceto, C. (2014). Displaying contextual information reduces the costs of imperfect decision automation in rapid retasking of ISR assets. *Human factors*, *56*(6), 1036-1049.

Russo, J. E. (1977). The value of unit price information. *Journal of Marketing Research*, 14, 193–201.

Russo, J. E., Carlson, K. A., & Meloy, M. G. (2006). Choosing an inferior alternative. Psychological Science, 17(10), 899-904.

Sadler, G., Battiste, H., Ho, N., Hoffmann, L., Johnson, W., Shively, R., ... & Smith, D. (2016, September). Effects of transparency on pilot trust and agreement in the autonomous constrained flight planner. In *Digital Avionics Systems Conference (DASC), 2016 IEEE/AIAA 35th* (pp. 1-9). IEEE.

Sanders, M. S., & McCormick, E. J. (1998). *Human factors in engineering and design* (Vol. 25, No. 2, pp. 153-153). Emerald Group Publishing Limited**.**

Schkade, D.A., & Kleinmuntz, D.N. (1994) Information displays and choice processes: Differential effects of organization, form, and sequence. *Organizational Behavior and Human Decision Processes, 57*, 319-337.

Selkowitz, A., Lakhmani, S., Chen, J. Y., & Boyce, M. (2015, September). The Effects of Agent Transparency on Human Interaction with an Autonomous Robotic Agent. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 59, No. 1, pp. 806-810). Sage CA: Los Angeles, CA: SAGE Publications.

Seong Y., & Bisantz, A. M. (2008). The impact of cognitive feedback on judgment performance and trust with decision aids. *International Journal of Industrial Ergonomics*, *38*, 608– 625.

Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: an effort-reduction framework. *Psychological bulletin*, *134*(2), 207.

Shah, A. K., & Oppenheimer, D. M. (2009). The path of least resistance using easy-to-access information. *Current Directions in Psychological Science*, *18*(4), 232-236.

Siegrist, M., Leins-Hess, R., & Keller, C. (2015). Which front-of-pack nutrition label is the most efficient one? The results of an eye-tracker study. *Food Quality and Preference*, *39*, 183–190.

Simon, H. A. (1990). Invariants of human behavior. Annual review of psychology, 41(1), 1-20.

Simon, H. A. (1955). A behavioral model of rational choice. *The quarterly journal of economics*, *69*(1), 99-118

Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research*, 16, 158–174

Simonson, I., Bettman, J. R., Kramer, T., & Payne, J. W. (2013). Directions for judgment and decision making research based on comparison selection: Reply to Arkes, Johnson, and Kardes. *Journal of Consumer Psychology*,*1*(23), 161-163.

Slovic, P. (1972). From Shakespeare to Simon: Speculation—and some evidence—about man's ability to process information, Oregon Research Institute Bulletin, 12, 1–29.

Stone, D. N., & Schkade, D. A. (1991). Numeric and linguistic information representation in multiattribute choice. *Organizational Behavior and Human Decision Processes*, 49(1), 42-59.

Stowers, K., Kasdaglis, N., Newton, O., Lakhmani, S., Wohleber, R., & Chen, J. (2016, September). Intelligent agent transparency: The design and evaluation of an interface to facilitate human and intelligent agent collaboration. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 60, No. 1, pp. 1706-1710). Sage CA: Los Angeles, CA: SAGE Publications.

Sun, Y., Li, S., & Bonini, N. (2010). Attribute salience in graphical representations affects evaluation. *Judgment and Decision Making*, *5*(3), 151-158.

Sun, Y., Li, S., Bonini, N., & Su, Y. (2012). Graph-Framing Effects in Decision Making. *Journal of Behavioral Decision Making*, 25(5), 491-501.

Svenson, O., Karlsson, G. (1986) Attractiveness of decision alternatives characterized by numerical and non-numerical information. *Scandinavian Journal of Psychology*, 27, *74-84.*

Szalma, J. L. (2014). On the application of motivation theory to human factors/ergonomics: Motivational design principles for human-technology interaction. *Human Factors, 56, 1453–1471.*

Tannenbaum, D., Fox, C. R., & Goldstein, N. J. (2013, November). Partitioning menu items to nudge single-item choice. In 34th Annual Meeting for the Society for Judgment and Decision Making.

Tannenbaum, D., Doctor, J. N., Persell, S. D., Friedberg, M. W., Meeker, D., Friesema, E. M., ... & Fox, C. R. (2015). Nudging physician prescription decisions by partitioning the order set: results of a vignette-based study. *Journal of General Internal Medicine*, 30(3), 298-304.

Thaler, R. H., & Sunstein, C. R. (2008). Nudge: improving decisions about health, wealth, and happiness. Penguin.

Thaler, R. H., Sunstein, C. R., & Balz, J. P. (2012). Choice architecture. In E. Shafir (Ed.), *The behavioral foundation of policy*, Princeton University Press.

Towal, R. B., Mormann, M., & Koch, C. (2013). Simultaneous modeling of visual saliency and value computation improves predictions of economic choice. Proceedings of the National Academy of Sciences, 110(40), E3858-E3867.

Tufte, E. R., & Robins, D. (1997). *Visual explanations* (p. 52). Cheshire, CT: Graphics.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185, 1124–1131.*

Van der Laan, L. N., Hooge, I. T., De Ridder, D. T., Viergever, M. A., & Smeets, P. A. (2015). Do you like what you see? The role of first fixation and total fixation duration in consumer choice. *Food Quality and Preference*, *39*, 46-55.

Vessey, I. (1991). Cognitive Fit: A Theory-Based Analysis of the Graphs Versus Tables.

Weber, S. J., & Cook, T. D. (1972). Subject effects in laboratory research: An examination of subject roles, demand characteristics, and valid inference. *Psychological Bulletin*, *77*(4), 273.

Wickens, C. D., & Carswell, C. M. (1995). The proximity compatibility principle: its psychological foundation and relevance to display design. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *37*(3), 473-494.

Wickens, C. D., & Hollands, J. G. (2000). *Engineering psychology and human performance.* New York, NY: Prentice Hall.

Wickens, C., McCarley, J., & Steelman-Allen, K. (2009, October). NT-SEEV: A model of attention capture and noticing on the flight deck. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 53, No. 12, pp. 769-773).

Wolfe, J. M. (2007). Guided search 4.0. *Integrated models of cognitive systems*, 99-119.

Wong, J. (2012). The compromise effect in market settings. *Psychology Honors Papers*, Paper 21.

Wright, J. L., Chen, J. Y., Barnes, M. J., & Hancock, P. A. (2016a). The Effect of Agent Reasoning Transparency on Automation Bias: An Analysis of Response Performance. In *International Conference on Virtual, Augmented and Mixed Reality* (pp. 465-477). Springer, Cham.

Wright, J. L., Chen, J. Y., Barnes, M. J., & Hancock, P. A. (2016b). Agent Reasoning Transparency's Effect on Operator Workload. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 60, No. 1, pp. 249-253). Sage CA: Los Angeles, CA: SAGE Publications.