Abstract

People often must make inferences in domains with limited information. In such cases, they can leverage their knowledge from other domains to make these inferences. This knowledge transfer process is quite common, but what are the underlying mechanisms that allow us to accomplish it? Analogical reasoning may be one such mechanism. This dissertation explores the role of analogy in influencing decision-making performance when faced with a new domain. We delve into the knowledge transferred between tasks and how this influences decision-making in novel tasks. Experiment I has two conditions, and each condition has two tasks. In one condition, the two task domains are analogically related, where for example, participants make inferences first about water flow and then about heat flow. In the second condition, the domains do not share obvious similarities. For example, car efficiency and water flow. Experiment I shows that participants presented with an analogy demonstrated better performance than those without. We hypothesize that this knowledge transfer occurs in two ways: firstly, analogical mapping enhances comprehension of cue utilization in a new task; secondly, the strategy employed is transferred. In Chapter 3, we developed a machine learning technique to uncover the strategies used by participants. Our findings reveal that the best-performing strategy from the old task is typically carried over to the new task. In Chapter 4, we developed a model of analogical transfer in multi-attribute decision making. We use the ACT-R theory of cognition as a framework to model knowledge transfer by integrating a reinforcement learning model of strategy selection with a model of analogy. The simulation results showcase a similar trend of both accuracy and strategy use to the behavioral data. Finally, we critically analyze our study's limitations and outline promising directions for future research, thereby paving the way for a deeper understanding of knowledge transfer mechanisms.

Analogical Transfer in Multi-Attribute Decision Making

by

Jun Fang

B.S., Iowa State University 2012
M.S., Iowa State University 2015
M.S., Syracuse University, 2017

Dissertation
Submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Cognitive Psychology.

Syracuse University
June 2023

Acknowledgments

It's hard to fathom that my journey at Syracuse University is drawing to a close. The road to this point has been far from smooth, strewn with trials and tribulations, yet it has given me knowledge and life experiences that I could not have acquired elsewhere. As I pen these words, my heart brims with gratitude.

First and foremost, I want to express my profound gratitude to my Ph.D. advisor, Lael Schooler. His guidance, consistent encouragement, and unwavering faith in my abilities have been invaluable. As an exemplary scholar—intelligent, open-minded, humble, and dedicated—he has been a true role model. He has provided me immense support in my studies while also giving me the freedom to explore my passions and interests. Thanks to him, I have learned ACT-R modeling and connected with colleagues in the community. Lael's understanding nature during my difficult times and his never-ending patience with periods of low productivity have proven him a genuinely compassionate individual. He taught me to face challenges with calm and courage, a lesson that will serve me well throughout my life.

My sincere appreciation goes to the Applied Statistics program in the Math department, which I discovered early in my Ph.D. journey. I recall how the first spark of a research idea came from the Machine Learning course I took with Prof. Wang Yi. I also cherished Prof. Chen Pinyuen's class on ranking and selection, and the challenging but incredibly rewarding course with Prof. Kim. These courses provided a robust foundation for my research and led me to make lasting friendships. I hold dear the memories of Thanksgiving dinners at Professor Chen's house and the camaraderie we shared.

I found a home in the Department of Psychology's CBB program at Syracuse. My heartfelt thanks go to Professor Mike Kalish and Professor David Kellen, who both served on my Master's thesis and qualification exam committees. Without their constructive feedback and thought-provoking questions, I wouldn't have reached this pivotal point of defending my dissertation.

Lastly, I wish to express my deep gratitude to my husband, Yu Chaoran, for his unending support and technical troubleshooting. To our parents, whose unconditional love has been my most formidable support throughout this journey, I am eternally thankful. Your strength continues to inspire me.

Table of Contents

List of Tables

List of Figures

Chapter 1: Introduction

In daily life, we frequently face tasks in ever-changing environments, necessitating predictions or decisions in unfamiliar domains. Rather than starting from scratch, we rely on prior knowledge and experience to navigate new situations, a process called knowledge transfer. This phenomenon, observed in individuals as young as 3 years old (Brown & Kane, 1988), is essential for rapid adaptation to novel situations with limited information. Knowledge transfer allows us to apply existing expertise and skills to make better decisions and predictions in unfamiliar contexts. This concept is also central to machine learning research, as automated systems need to transfer knowledge between domains for rapid learning (Pan & Yang, 2010). Thorndike was the first psychologist to systematically study transfer (Thorndike & Woodworth, 1901). They posited that the mind consists of specific habits and associations, not general faculties, providing a range of narrow responses to particular stimuli. Thorndike's theory of identical elements states that training in one activity transfers to another only if they share common stimulus-response elements. However, this view does not account for intelligent adaptation or reconstruction.

Singley and Anderson (1989) expanded on Thorndike's idea, suggesting that transfer between individual skills occurs only if they share identical production rules—specific action sequences involved in performing a task. Thus, effective knowledge transfer requires a deep understanding of individual skills and awareness of their commonalities and differences. Taatgen (2013) further refined this concept, breaking production rules into reusable elementary building blocks for other tasks and skills. Determining the possibility of transfer between tasks can be complex and nuanced, often dependent on identifying shared elements which may relate to cognitive processes, knowledge types, or specific strategies. Ultimately, successful knowledge transfer requires a thorough analysis of tasks and an in-depth understanding of relevant cognitive processes and principles.

Analogy can be the conduit for knowledge transfer that involves establishing connections between old and new environments by identifying and correlating similarities in the relationships among concepts. Seminal research by Gentner (1983), Falkenhainer, Forbus, and Gentner (1989), Holyoak and Thagard (1989), and Hummel and Holyoak (1997) has extensively studied this concept and offered valuable insights into its workings. While analogies can be a powerful cognitive tool, they are often insufficient and require integration with other cognitive tasks. This has led to increased interest in combining analogy models with various cognitive tasks (Forbus et al., 2017).

> Exploring the roles of analogy and similarity in a broader range of cognitive processes via large-scale simulation is an enterprise that is just a beginning. By making available a robust model of analogical matching, we hope we can encourage others to join us in these investigations. (p.1193)

In this dissertation, my aim is to delve into the decision-making processes that utilize analogical transfer, a mechanism that enables the extraction of relevant information based on previously acquired knowledge exhibiting structural similarities. My hypothesis posits that leveraging knowledge from prior tasks can impact the precision of information selection and decision-making in complex tasks involving multiple attributes.

This inquiry is situated within the framework of ecological rationality (Goldstein and Gigerenzer, 2002). In this context, heuristics are not simply viewed as imperfect versions of intricate optimal statistical procedures that are presumed to be beyond the capacity of ordinary minds. Rather, they are regarded as adaptive strategies that have evolved synergistically with fundamental psychological mechanisms.

The merit or effectiveness of a heuristic strategy is not an absolute measure but rather is contingent on its specific adaptability to the environment in question. In this light, analogies serve as a heuristic tool. They empower decision-makers to isolate a crucial subset of information from a larger body of data, thus facilitating prompt and precise decisions. This reflects the essence of ecological rationality: cognitive processes, like the use of analogies as heuristics, are assessed based on their alignment and adaptability within a given environmental context.

My objective is to examine the efficacy of analogical transfer in improving decision-making performance and delve into the underpinning mechanisms and conditions that promote such transfers. Given the existing research gap in the intersection of analogy and decision-making studies, I aim to contribute to a deeper understanding of knowledge transfer in the context of decision-making.

Analogy and decision-making are intrinsically connected in numerous tasks, especially in lab experiments where decision-makers are presented with a cover story to support their decision-making process. The cover story generally includes relevant information that enables decision-makers to comprehend the importance and direction of the task's attributes. For instance, a decision task that requires choosing the most profitable company from two contenders. Factors such as the companies' expense ratios, market shares, and other elements influencing their profitability are considered (e.g., Rieskamp & Otto, 2006; Rieskamp & Hoffrage, 2008). The cover story elucidates the relationships among these factors, acting as a conduit to draw upon analogs from previous knowledge. Despite this, the discourse bridging decision-making and analogy is still nascent. While evident, the interplay between these two areas requires further exploration and understanding.

Analogies prove advantageous when previous knowledge is structurally similar to the concepts present in the current task (e.g.,Gaventti et al., 2005; Gentner et al., 2003; Gentner, 2017; Forbus et al., 2017; Holyoak, 2012; Richland et al., 2015; Salvucci & Anderson, 2001). Decision-makers can tap into their past experiences to make well-informed decisions, particularly when they discern similarities between the current task and their previous experiences. For instance, if a decision-maker possesses expertise in the retail industry, they may leverage this knowledge to assess the profitability of the two companies involved in the current decision task. This approach demonstrates the integral role of analogical thinking in facilitating effective decision-making across various contexts. However, what is the mechanism by which analogies aid decision-making? To answer this question, one place to explore is the heuristic decision strategies in multi-attribute decision-making.

Strategy Selection

Decision-makers must not only understand the attributes of tasks but also select appropriate strategies for making decisions. Researchers have proposed that the mind possesses an "adaptive toolbox" comprising a collection of heuristics (Gigerenzer, Todd, & the ABC Research Group, 1999; Todd & Gigerenzer, 2000; Marewski et al., 2010). Many heuristics can be classified based on their reliance on knowledge and the accessibility of memories. Knowledge pertains to the cues available to decision-makers, while memory-based heuristics require decision-makers to search their memory for useful cues to inform their decisions.

To illustrate heuristic strategies, we'll employ a classic example from analogy literature—the water flow system (Falkenhainer et al. 1989)—as it serves as our experimental material, which we will revisit later. Water flows from the large beaker to the small vial through a pipe, as shown in Figure 1-1. The water flows because the pressure of the large beaker is greater than that of the small vial. The information about the two water flow systems is illustrated in a table shown in Figure 1-2, which presents five cues ranked by their validities. A cue's validity is defined as the probability of making a correct inference when the cue discriminates (i.e., when the cue values differ for the two alternatives being compared). In a paired-comparison task, the goal is to infer which pair alternative has a larger criterion value. In this case, the task is choosing the water flow system with the higher flow rate.

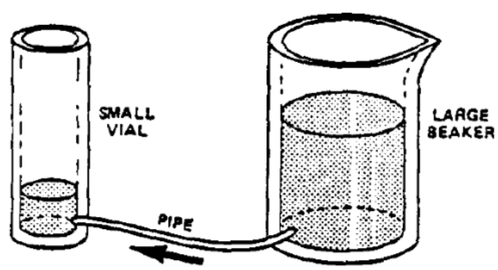

Figure 1-1. The water flow system. The picture is updated from Falkenhainer et al. (1989).

Knowledge-based heuristics employ various cues to make inferences. For instance, the take-the-best heuristic (Gigerenzer & Goldstein, 1996) considers cues in the order of their validities. When a cue discriminates, the decision maker selects the alternative with the cue value corresponding to a higher criterion value. If the cue values of the alternatives are identical, the decision maker proceeds to the next cue. If no cue discriminates, a random alternative is chosen. Table 1-1 demonstrates how take-the-best and three other knowledge-based strategies would approach the water flow system decision shown in Figure 1-2. In order to track decision makers' decision processes (e.g., information search order, mouse trajectories), we translate the visual representations, such as Figure 1-1, to the abstract representations in Figure 1-2, for them to use.

	Cue validity	*System A*	*System B*
Water flow rate		low	high
Pipe diameter	0.9	low	medium
Pipe length	0.76	high	medium
Water pressure difference	0.60	low	high
Beaker diameter	0.55	high	low
Via diameter	0.5	medium	medium

Figure 1-2 The water flow system is described by five cues. Water flow rate is the criterion variable whose values are unknown to the decision maker. The decision maker needs to infer from the cues which system has a higher water flow rate. The cues are ordered by validity, a measure of a cue's quality. It takes on three values of low, medium, and high.

Δ-inference operates similarly to take-the-best (Luan et al., 2014); however, it differs in that it stops searching and makes a decision when the cue value of one alternative surpasses that of the other by a threshold Δ. Take-the-best can be regarded as a special case of Δ-inference when Δ is set at zero for all cues. For both take-the-best and Δ-inference, we expect a cue-wise information search, meaning decision makers would examine both alternatives' values for a cue before moving on to the next cue or making a decision.

Another strategy, weighted-additive (WADD), involves weighing an alternative's cue values by each cue's importance, summing the weighted values for an overall score, and choosing the alternative with the highest score (Payne et al., 1988). In a task with binary cues, cue values can be weighted by each cue's validity. When cues are not binary, cue dichotomization may simplify the weighting-and-adding process. Specifically, one may first dichotomize a cue with a threshold, treating the higher or more favorable values as "1" and others as "0". A weighted score for each alternative is then calculated based on the dichotomized cue values and cue validities. In the current task, a set of trichotomized cues (low, medium, high) are used. One way of dichotomization is to assign 1 to “high” and the rest to 0.

Tallying is a special case of WADD in which all cues are considered equally important (Payne et al., 1988). This approach can significantly reduce computational demands. For WADD and Tallying, an alternative-wise search is used (Canellas & Feighm, 2017; Rieskamp & Hoffrage, 2008), meaning one would inspect all cue values of one alternative before examining those of the other alternative.

These four strategies can be classified into two general categories based on their use of cue information. Take-the-best and Δ-inference exemplify noncompensatory strategies, wherein favorable or unfavorable values on lower-ranked cues cannot compensate for values on higher-ranked cues and thus cannot override decisions made by higher-ranked cues. For example, one might insist on purchasing a four-wheel-drive Jeep regardless of the dealer's discount on a two-wheel-drive model. Conversely, WADD and Tallying represent compensatory strategies that allow trade-offs among cues, enabling favorable values on lower-ranked cues to compensate for unfavorable values on higher-ranked cues.

Strategies	Examples
Take-the-best	The pipe diameter cue is the most valid cue, followed by the other cues in order. When a cue discriminates (i.e., the two systems have different values on the cue), a decision maker decides by using this cue; otherwise, they look up the next cue. For example, System A has a value of "low" and System B has "medium". TTB chooses A.
Δ-inference	The pipe diameter cue only discriminates between the two water systems when their values differ by more than a threshold Δ. When the Δ for pipe diameter is set at 2 units, it will not discriminate in this case, because the difference between the two cars is only 1 unit (i.e., low vs medium). A person should then move on to the next cue. The decision should be made on the third cue "Water pressure difference" because the difference in cue values is greater than 1 unit. Δ-inference chooses B.
Weighted-additive (WADD)	A decision maker first evaluates each water system's value on a cue, translating it to a binary value by coding favorable values as 1 and unfavorable values as 0. For instance, system A's pipe diameter is low, which is small and thus coded as 0. The dichotomization threshold can vary by person and task. In the example, we consider "high" as 1 and the rest are "0". After dichotomization, the person derives an overall score for each car by multiplying the binary cue values with the corresponding cue validities and chooses the system with a higher score. For example, the overall score for System A would be: 0.9 x 0 + 0.76 x 1 + 0.60 x 0 + 0.55 x 1 + 0.5 x 0 = 1.31.
Tallying	A decision maker counts the number of cues in favor of one system and chooses the system with more favorable cues. For example, system B has a score of 1, because it has higher water pressure, while system A has a score of 2 because it has a longer pipe length and larger beaker diameter. Thus, the person chooses system A.

Table 1-1 Examples of Strategy Applications for the Decision Problem are Shown in Figure 1-2

Accessibility-based strategies rely on the ease with which information can be retrieved. The fluency heuristic, for instance, posits that items perceived as more quickly recognized are considered to have a larger criterion value than those perceived as more slowly recognized. Another example is the recognition heuristic, which suggests that recognized objects have a larger criterion value than unrecognized ones. In this study, however, we primarily focus on knowledge-based heuristics.

A central question in the simple heuristics research program is understanding when and how people utilize different heuristics (Gigerenzer, Todd, & the ABC Research Group, 1999; Todd & Gigerenzer, 2000; see Marewski, Gaissmaier, & Gigerenzer, 2010 for an overview). Two primary approaches explain strategy selection. One approach is tied to learning, known as the cost-benefit approach. This theory proposes that individuals balance the costs (such as time and effort) against the benefits of using a particular strategy (Beach & Mitchell, 1978; Christensen-Szalanski, 1978; Payne et al., 1988, 1993). The other approach focuses on selection from available resources, called the cognitive niche approach. This concept suggests that the choice of strategy hinges on the interaction between fundamental cognitive capabilities and the structure of the environment (Marewski & Schooler, 2011).

Many quantitative theories of strategy selection within the cost-benefit framework emphasize the role of learning (Busemeyer & Myung, 1992; Erev & Roth, 2001; Rieskamp & Otto, 2006). For instance, Rieskamp and Otto's (2006) model proposes that people choose among different strategies by learning from their feedback. Similarly, reinforcement mechanisms in the ACT–R, a unified theory of cognition, favor selecting cognitive processes that implement successful strategies (Fu & Anderson, 2006; Lovett & Anderson, 1996). ACT-R assumes that cognition can be understood in terms of a set of basic principles that govern the operation of a set of specialized cognitive modules (e.g., declarative memory, perception), which interact through a centralized production system (Anderson et al., 2004). Numerous studies have demonstrated that individuals learn to adapt to various inference problems and can make optimal inferences (e.g., Anderson, 1991; Ashby & Maddox, 1992).

In this dissertation, we concentrate on the cost-benefit approach to strategy selection, where individuals select strategies based on the feedback they receive from making decisions. Participants repeatedly engaged in inferential tasks under diverse environmental conditions. Sometimes, these tasks exhibited analogical relationships, particularly when accompanied by outcome feedback. The objectives of the behavioral experiments are twofold: (1) to investigate the potential for participants to use an analogy to transfer knowledge from an old environment to a new task, and (2) to replicate the findings of prior studies showing that participants can adapt to different environments by learning the best-performing strategy.

Dissertation Outline

This dissertation investigates how analogy connecting two environments affects knowledge transfer in multi-attribute decision-making.

In Chapter 1, an introduction to the research problem at hand is provided.

Chapter 2 investigates analogical transfer's impact on multi-attribute decision-making through a behavioral experiment. Participants were directed to make decisions in two distinct settings. The first task served as a training platform for the participants, aiding in understanding cues and optimizing performance. Specifically, half of the participants were provided with an analogous scenario to their subsequent task, while the remainder were exposed to an unrelated task. In the second task, all participants were given the same task, solely with attribute names with no attribute importance. The group trained on the analogous task was predicted to apply cue validities and strategies to the second task. The objective here is to compare the performances and strategic choices of the two groups across both tasks.

Chapter 3 explores how participants select strategies in the two environments via a machine learning strategy identification approach (Fang, Schooler, and Luan, 2022). The MLSI approach uses the collected behavioral data to train machine learning models, which can then identify the strategies participants selected to adapt to the environments. We first demonstrate the effectiveness of this machine learning approach by contrasting it with other strategy recovery methods, then apply this approach to the analogy experiment. By identifying the strategy selection processes in both environments, this chapter tests how the first environment affects strategy selection in the second environment.

Chapter 4 employs the ACT-R (Adaptive Control of Thought-Rational) cognitive architecture to simulate the strategy selection process within the two environments and the analogical reasoning process during the transition between them. ACT-R is a cognitive architecture that outlines the cognitive processes involved in perception, memory, and decision-making. Various heuristic strategies are implemented as a sequence of steps, each associated with a utility value. This value is updated with each feedback received by the simulated agent. Positive feedback boosts the strategy's utility, while negative feedback diminishes it. The likelihood of selecting a particular strategy is governed by its associated utility. This is the basis of learning a specific strategy in the selection process. The existing analogical mapping method is coupled with strategy selection to model behavior when decision-makers employ an analogy. This chapter aims to shed light on the cognitive mechanisms driving analogical transfer in decision-making using ACT-R. It addresses the call for integrating an analogy model with other cognitive tasks, thereby comprehensively examining these interlinked cognitive processes.

Chapter 5 encapsulates the findings and conclusions and deliberates on the study's limitations and potential future trajectories.

Overall, this dissertation seeks to contribute to understanding how an analogy can affect knowledge transfer in multi-attribute decision-making tasks. By exploring the different aspects of knowledge transfer, including strategy selection, cognitive processes, and knowledge representation, the dissertation aims to comprehensively understand how analogical reasoning can impact knowledge transfer in these tasks.

Chapter 2

Analogical Transfer in Multi-attribute Decision Making

Humans must make inferences with limited time and knowledge in many real-world scenarios. In the process of making decisions, individuals typically engage in an information search to gather relevant data. For example, someone looking to purchase a reliable used car must choose between two options, necessitating the collection of attributes or cues such as mileage, model year, and accident histories to make an informed decision (Stewart, 1988). Simple heuristics, such as take-the-best, demand only a few cues to reach a decision, making the selection of appropriate cues vital for decision accuracy. In this approach, cues are ranked by importance, with decision-makers comparing the values of the most critical cue and ending the search when values differ.

However, how do individuals determine which cues to prioritize? One theory suggests that people search their memory for cues, using the first cue they recall (take the first, e.g., Bröder and Gaissmaier, 2007; Gigerenzer and Gaissmaier, 2011; Marewski and Schooler, 2011) because the accessibility of memories is informative. Highly valuable cues receive more exposure in various environments (e.g., the press and the internet) and are more readily accessible in memory. The likelihood of retrieving an item correlates strongly with its presentation probability in these environments (Anderson and Schooler, 1991). Therefore, the first cue that comes to mind often has a high criterion value. Another theory posits that decision-makers search for and learn cues (Todd and Dieckmann, 2004), selecting appropriate cues in a learning-while-doing situation and receiving feedback on their decisions' adequacy. Thus, they learn cue validities through trial and error. In a simulation study, Todd and Dieckmann (2004) showed that it took over 100 decisions to find the ranking of nine natural cues that matched the true cue validities. Consequently, heuristic strategies that are easy to apply may not necessarily be easy to set up (Dougherty et al., 2008). But how do decision-makers navigate situations where they cannot retrieve direct information quickly and lack sufficient time to learn cue validities? This is where past knowledge becomes crucial.

In another scenario, leveraging past knowledge becomes beneficial for individuals as they navigate changing environments. While many studies focus on observing people making repeated decisions within a single domain with constant cues, research has demonstrated that individuals are sensitive to task changes and can adapt their information search strategies to accommodate varying task characteristics, such as time pressure (Rieskamp & Hoffrage, 2008; Bobadilla-Suarez et al., 2018), the number of alternatives to be evaluated (Katz et al., 2010; Van Ravenzwaaij et al., 2014), the cognitive cost associated with each strategy (Fechner et al., 2018; Marewski & Schooler, 2011), and environmental probabilities (Nelson et al., 2014). Even minor alterations in the environment or tasks can influence people's information search and strategy selection (Lee et al., 2014).

The advantage of past knowledge lies in enabling decision-makers to apply their prior experience and understand new situations effectively. For instance, a coffee enthusiast visiting Asia, where tea is more prevalent, can adeptly adapt by familiarizing themselves with various types of tea and their preparation methods, drawing upon their prior knowledge of coffee. Both coffee and tea share similarities as plant-derived beverages, though they possess distinct characteristics. Brewing methods for coffee, such as drip, pour-over, or espresso, have parallels in tea preparation, including steeping in a teapot or using a tea infuser. Furthermore, the strength spectrum of coffee, from light to dark roast, mirrors that of tea, with delicate green teas to robust black teas. Additionally, both beverages contain caffeine, with varying amounts across different types. Individuals transfer their prior knowledge to navigate new environments or domains, mapping concepts from familiar to unfamiliar ones.

To address the challenge of selecting the most suitable cues and strategies at the beginning of a task, we propose leveraging analogical transfer, which involves ranking cues based on previously learned structurally similar knowledge. This process of transferring knowledge from prior tasks can enhance cue selection and strategy selection in multi-attribute decision-making tasks, allowing individuals to navigate new situations more effectively.

Knowledge Transfer

One way to adapt to a new environment is to compare it with a familiar one and identify what is similar and what is different. These similarities and differences can involve specific examples, conditions, characteristics, or patterns of relationships. Analogies are a type of similarity that focus on the patterns of relationships. Gonzalez et al. (2003) proposed an instance-based learning theory (IBLT) that explains how people transfer instances from memory to new situations when making decisions. According to this theory, people store instances as attributes and associated outcomes in their memory and retrieve them when they face a similar situation. Alternatively, Canini et al. (2010) model transfer learning via parameters. Parameter transfer involves constructing a prior distribution over probabilistically dependent categories, where knowledge about one category influences the distribution over others. Lastly, transfer through structural relations can be found in analogies. This approach has been extensively examined, with notable contributions from Gentner (1983), Falkenhainer, Forbus, and Gentner (1989), Holyoak and Thagard (1989), and Hummel and Holyoak (1997).

Analogical transfer, a form of structural knowledge transfer, is a fundamental and pervasive aspect of human cognition. This process involves identifying and utilizing correspondences between concepts or cues, with the core principle being that an analogy maps knowledge from one environment to another. This transfer conveys a system of relations known to hold in the familiar environment and potentially applicable in the new one. For example, Falkenhainer et al. (1989) studied water flow and heat flow, as shown in Figure 2-1. In their study of the analogy between water flow and heat flow, Falkenhainer et al. (1989) demonstrated how water flows from a large beaker to a small vial through a pipe due to the difference in pressure between the two containers. This is similar to how heat flows from a warm cup of coffee to an ice cube because the temperature of the coffee is higher than that of the ice cube. In this analogy, the concepts of pressure and temperature in the two domains can be mapped to each other through structural alignment. The structural alignment between two domains is achieved by highlighting the structural relations (Gentner, 1983).

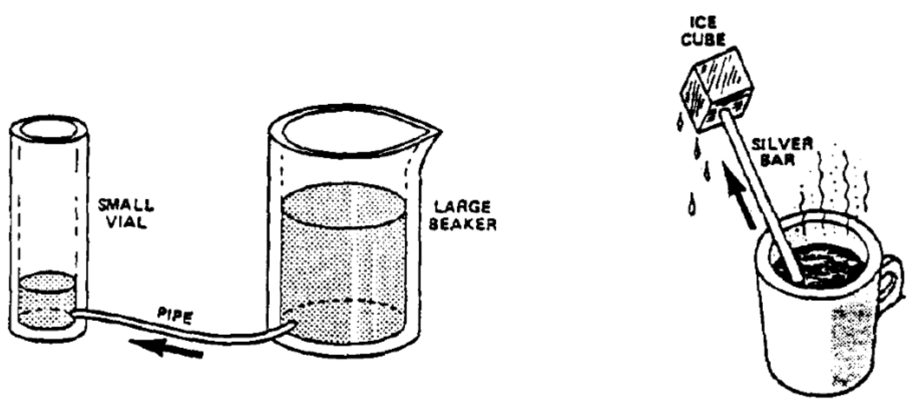

Figure 2-1. Two analogical physical situations: Water flow and heat flow (Falkenhainer et al. 1989)

According to the Structural Mapping Engine (SME) theory of analogy by Falkenhainer et al. (1989), a domain's knowledge representation is composed of first- and second-order logic. First-order logic consists of sentences with a single predicate, like "the pressure in the beaker is greater than the pressure in the vial." Second-order logic entails sentences that connect two first-order logic sentences, for instance, "the large beaker's pressure is greater than the vial's, which causes water to flow from the beaker to the vial." Knowledge representations for water flow and heat flow can be seen in Figure 2-2.

SME adheres to the systematicity principle: "A predicate that is part of a mappable system of interconnected relationships is more likely to be imported into the target domain than an isolated predicate." Therefore, alignment typically starts with predicates in second-order logic sentences, such as "cause" and "greater," and proceeds to map the attributes in first-order logic. In the example provided, both domains exhibit a causal relationship and a comparative scenario, leading to the mapping of water pressure to coffee temperature and the pipe to the silver bar. Other mappings and directions are also plausible in this analogy. The optimal mapping result is chosen based on scores derived from the number of matched predicates and the depth of the hierarchical structure among all possible mapping outcomes. In the context of multi-attribute decision-making, concepts like the pipe in the water system and the bar in the heat system act as cues within task environments.

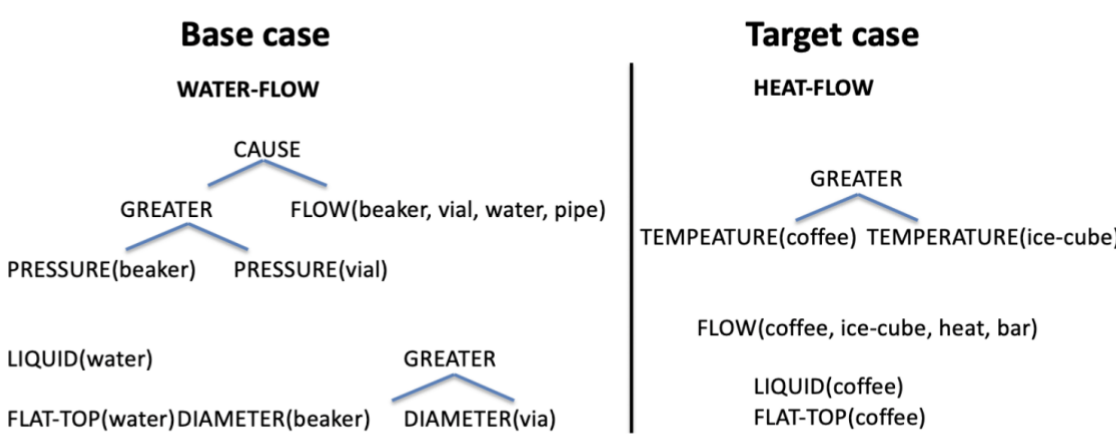

Figure 2-2. Hierarchical structures of water flow and heat flow. Adapted from Falkenhainer et al., (1989)

In this study, we aim to address several research questions related to decision-making processes and the role of analogical transfer. Specifically, we seek to understand how individuals prioritize cues when employing simple heuristics and how past knowledge influences cue selection and decision-making accuracy in new or unfamiliar domains. Furthermore, we aim to investigate whether analogical transfer can enhance cue selection and decision-making accuracy in multi-attribute decision-making tasks and identify the underlying mechanisms and conditions that facilitate analogical transfer in these contexts. Lastly, we explore how individuals adapt their information search and strategy selection in response to changing environments or tasks with varying characteristics.

Experiment I

Our study explored whether providing performance feedback could improve decision-making in repeated inferences and whether analogical transfer could improve decision accuracy. Participants made decisions in a non-compensatory environment, where the take-the-best strategy yielded the best results. They received feedback for each trial and earned one point for accurate decisions. The first part of our experiment served as a conceptual replication of Rieskamp and Otto's (2006) Study 1. According to their findings, participants learn the most effective strategy through feedback received in both compensatory, where strategies like tally and WADD are preferred, and non-compensatory environments, where non-compensatory strategies are preferred. In our study, we provided feedback on the take-the-best strategy, emphasizing the importance of cue validities. Once participants had mastered the optimal strategy in the noncompensatory environment, they proceeded to a second task where cue validities were not explicitly provided. They could apply the knowledge they acquired (e.g., cue validities) from the initial task to inform their decisions in the subsequent task. The two environments were connected via analogy. Our main goal was to investigate whether analogical transfer from previous experiences could impact decision-making in the current task.

Method

Participants and Design.

One hundred and twenty participants recruited from the SONA subject pool were randomly assigned to one of four between-subject conditions, as shown in Table 2-1. The participants assigned to each condition went through two tasks. Two participants in the water-to-heat condition and three participants in the car-to-heat condition were excluded due to the incompletion of the task. That resulted in 116 participants in total.

There were two analogy conditions and two no-analogy conditions. First, the base task where the cue validities were shown to them, and they learned how to use the cue validities and cue values to make decisions. Next, participants were given a new task and a set of cues without cue validities in the target task. In this case, participants did not know how important each cue was and hence could not apply strategies relying on cue ranking immediately, such as take-the-best. We use the water flow and heat flow analogy, as shown in Figure 2-1, for the analogy conditions. Before participants proceeded to the decision-making tasks, they were given a cover story of either how the water flow system works or how the heat flow system works. The script of the water flow system was "Water flows from the large beaker to the small vial through a pipe, as shown in the picture below. The water flows because the pressure of the large beaker is greater than the pressure of the small vial." After explaining how the water flows between two containers, participants introduced the five cues that can affect how fast the water flows. They learned the meaning and the direction of cues one at a time. The script of the heat flow system was “Heat flows from the warm coffee to the ice cube through a silver bar, as shown in the picture below. The heat flows because the temperature of the coffee is greater than the temperature of the ice.” Participants were then introduced to the five cues affecting how fast the heat flows between coffee and ice. We use the car efficiency task for the non-analogy condition. In the car efficiency task, participants were told to select one car among two cars with the highest gas efficiency given five cues, such as car price, mileage, age, number of accidents, and maintenance frequency. The car efficiency task was always the first, followed by either the water flow or heat flow task.

Upon comprehending the cover story, participants initially entered the base task. In this task, they were instructed to choose between two water (or heat) systems based on which one had a higher flow rate for water (or heat). Each system was characterized by five cues along with their respective validities. The water flow system task featured cues such as pipe diameter, pipe length, water pressure difference, beaker diameter, and vial diameter, while the heat flow system task included bar diameter, bar length, temperature difference, cup diameter, and ice cube diameter. Each cue had three potential values, such as low, medium, and high. The two alternatives (systems) were presented on a computerized information board in a matrix format, with five rows and two columns, displaying cue values hidden behind boxes. Cue validities were displayed in parentheses next to the cue names (see Figure 2-3). Participants accessed cue values by clicking on the corresponding box and indicated their choices by selecting a button at the bottom of the screen. The experimental program recorded mouse locations with 5-ms precision.

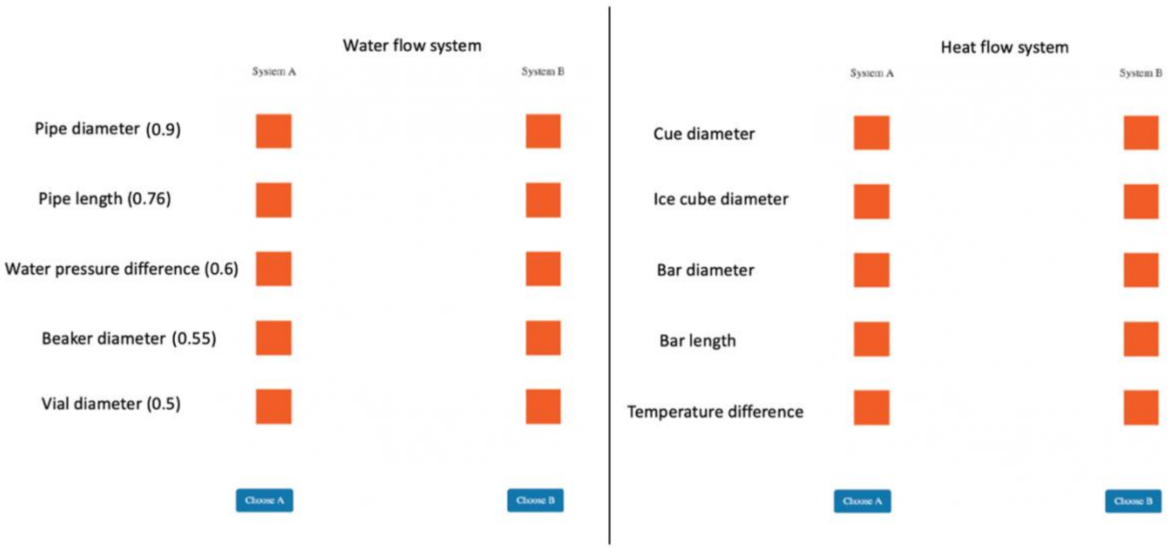

Figure 2-3 The experiment interface in the analogy condition where the base task is the water flow system, and the target task is the heat flow system. The figure represents two different screens in two conditions. In the water flow system task (base), the water flow systems were described by five cues, and the cue validities were shown in the parenthesis. The cue order was ranked by their validities. The cues were randomized in the heat flow system task (target), and the cue validities were not given.

Participants in each condition needed to finish two tasks. In both tasks, participants made 100 choices under no time pressure. The first 20 trials were pre-test, meaning no feedback was given, followed by four trial blocks, each containing the same set of 20 items. The first block measured the prior knowledge of the task and the potential initial preferences for one of the strategies. For the following four blocks, we provided participants with the outcome feedback to allow for learning. The base task and the target task were counterbalanced for the analogy conditions. The car's gas efficiency was unrelated to any water and heat tasks. We used this task in the control conditions to compare the learning results with the analogy conditions. The water flow and heat flow tasks are connected by an analogy, where the cues in the base task are expected to be mapped onto the cues in the target task. In the control task, participants were asked to choose which car had higher efficiency (measured in miles per gallon) between two cars given five relevant cues (model year, mileage, price, number of accidents, car maintenance). This is the control task because it is not analogous to water flow or heat flow tasks.

Conditions	Pre-test (no feedback)	Base task (with cue validity)	Pre-test (no feedback)	Target task (no cue validity)
Analogy condition		Water flow		Heat flow
Analogy condition		Heat flow		Water flow
None analogy		Car's gas efficiency		Heat flow
None analogy		Car's gas efficiency		Water flow

Table 2-1 The four conditions in the experiment. Each condition contains two tasks, the base task, and the target task. Initial knowledge was assessed in the pre-test trials with no choice outcome feedback. The feedback was then given for the rest of the task.

As described previously, we assume that participants choose from three candidate strategies, take-the-best, Δ-inference, and tally. For all trials, the item set was constructed such that take-the-best reached an accuracy of 0.9 (i.e., 18 correct of a possible 20 predictions), and Δ-inference's accuracy was 0.8 (i.e., 16 correct out of 20 predictions). The compensatory strategy, tally (0.55 accuracy), was discouraged based on the feedback. This study operates within the context of the adaptive toolbox framework, which often studies these three strategies. We acknowledge the existence of the adjustable spanner framework, which postulates one comprehensive adaptive strategy (Krefeld-Schwalb et al., 2019; Newell et al.,2007). However, exploring this model falls outside the purview of our current investigation. A future direction of research may include examining how the involvement of analogies could influence parameters within the adjustable spanner model.

We used a training task after the base and target tasks to motivate participants to learn from the feedback. Participants were told they would learn the best strategy to use in the training phase. The total number of training trials is 180, but if they make one correct decision in the previous two tasks, the number of training trials will be reduced by one. To illustrate how the reward system works, we show Figure 2-4 in the instructions to the participants.

Figure 2-4 The rewarding system of the experiment.

For example, if a participant makes 80 correct decisions in the first task and 85 correct decisions in the second task, then the total number of training trials they will go through is 15. The procedure motivates participants to do well in the main tasks. The data collected in the training trials were used to train machine learning models that can identify the strategies they used throughout the experiment. We first show the data analysis results on the learning behavior using hierarchical latent Bayesian analysis. In the following chapter, we will discuss in detail how machine learning techniques come into play in identifying the strategies used.

Behavioral Analysis

Participants learned to make decisions across the ten blocks for all four conditions in the task environment. We show each block's performance in terms of decision accuracy of each block in Figure 2-5.

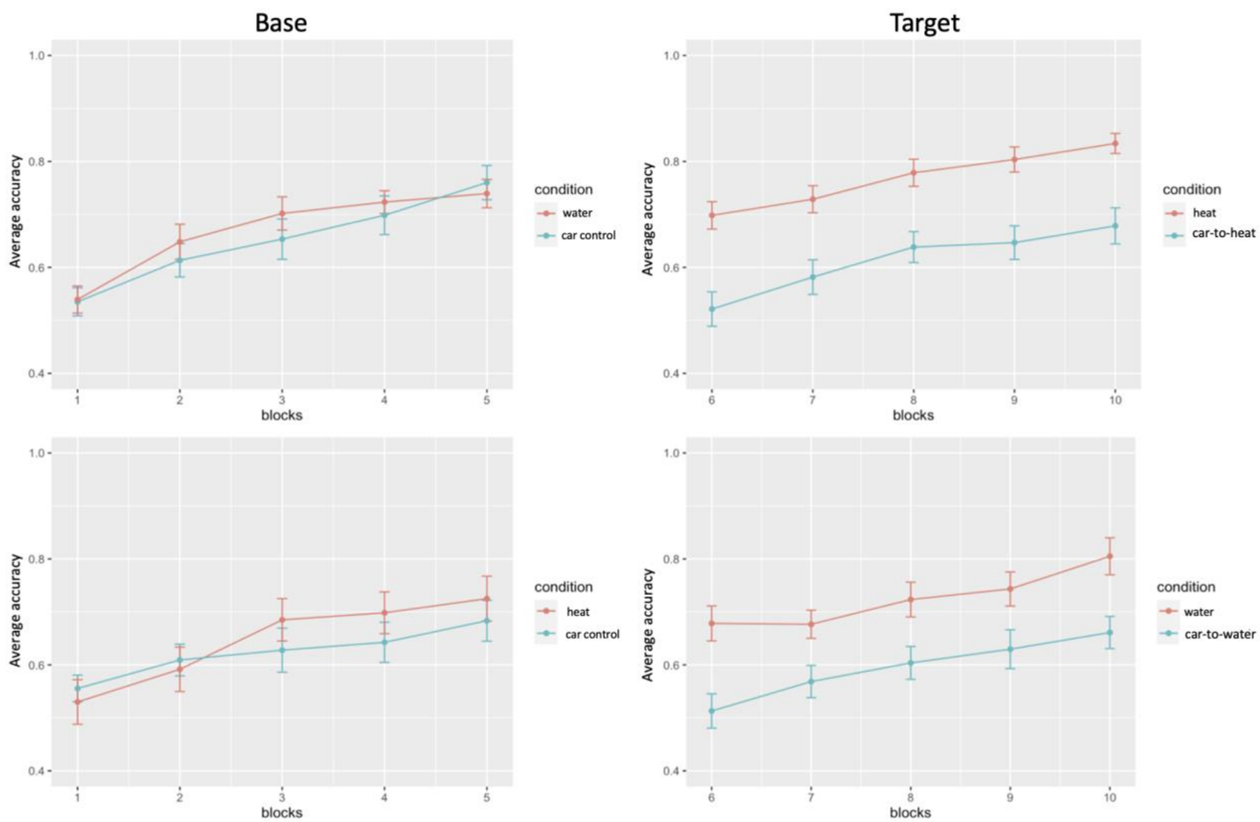

Figure 2-5 The average accuracy rate of the participants across the five trial blocks for four experiment conditions. The upper two graphs are the learning curve for the water flow (the blue line in the base task) to the heat flow (the blue line in the target task) and the learning curve for the car efficiency task (the red line) to the heat flow task (The red line). The lower graphs show the learning curve for the heat flow to the water flow task (the blue lines) and the learning curve for the car efficiency task to the water flow task (the red lines).

We report Bayes factors for all conditions and blocks. The Bayes factors were estimated using JASP (JASP Team, 2017; Morey & Rouder, 2015). We use the notation BF_{10} to indicate evidence that favors the alternative hypothesis. In other words, if the Bayes factor is between 1 and 3, it is weak evidence for the given alternative hypothesis. If the Bayes factor is between 3 and 20, it is positive evidence. It is strong evidence that the Bayes factor is between 20 and 100. It is very strong evidence if the Bayes factor is greater than 100. The evidence favors the null hypothesis if the Bayes factor is less than 1 (Raftery, 1995). Table 2-2 shows ANOVA summaries and Bayes factors for the five blocks in the Base task. The performance of five blocks across the water flow, heat flow, and car efficiency tasks are not significantly different. Note that the two no-analogy group (car-to-water, car-to-heat) is aggregated as one group in the base task.

Base Task	Sum of Squares	df	Mean Square	F	p	BF_{10}
Block 1	0.004	2	0.002	0.077	0.926	1
Residuals	3.098	112	0.028			
Block 2	0.048	2	0.024	0.696	0.501	1
Residual	3.858	112	0.034			
Block 3	0.081	2	0.041	0.986	0.376	1
Residual	4.624	112	0.041			
Block 4	0.051	2	0.026	0.736	0.481	1
Residual	3.910	112	0.035			
Block 5	0.005	2	0.002	0.066	0.936	1
Residual	4.156	112	0.037			

Table 2-2 ANOVA Summary Table for the Base Task accuracy rates for water flow, heat flow, and car conditions.

When entering the target task, participants who previously learned about the water flow system task started with a higher accuracy rate in the first block than those who learned from the car efficiency task. They initiated the heat flow task with an accuracy rate of around 70%, similar to the other analogy group that studied heat flow first. The two no-analogy (i.e., car-to-water, car-to-heat) groups all started at a lower accuracy rate, with car-to-heat slightly higher but not significantly different from the other no-analogy group. Different from the base task, car-to-water and car-to-heat group were analyzed as two separate groups. The performance gap between the analogy and no analogy groups indicates the knowledge transferred from the base task to the current target task for the analogy groups. Table 2-3 shows the ANOVA summaries and Bayes factors for the five blocks in the Target task. The performance of five blocks across the water flow task, heat flow task, and car efficiency task are significantly different in the Target task, and the Bayes factors are consistent with the results of ANOVA.

Target Task	Sum of Squares	df	Mean Square	F	p	BF_{10}
Block 6	0.534	3	0.178	6.764	<0.001	84.338
Residuals	3.098	111	0.028			
Block 7	0.551	3	0.184	8.552	<0.001	589.582
Residual	2.385	111	0.021			
Block 8	0.628	3	0.209	10.470	<0.001	>1000
Residual	2.221	111	0.02			
Block 9	0.656	3	0.219	9.433	<0.001	>1000
Residual	2.573	111	0.023			
Block 10	1.075	3	0.358	12.251	<0.001	>1000
Residual	3.246	111	0.029			

Table 2-3. ANOVA Summary Table for the Target Task accuracy rates for water flow, heat flow, water-to-car, and heat-to-car conditions.

Growth Curve Model

We employ the Growth Curve Model (GCM) to analyze and compare learning trajectories in analogy and non-analogy conditions, as our data is longitudinal. Longitudinal data consists of T measurements for individual i (i = 1, ..., N) taken on occasions t (t = 1, ..., T), where N signifies the total number of participants in the sample. A primary advantage of the longitudinal design is its ability to offer insights into within-person change mechanisms or intraindividual variability. This approach enables us to identify person-specific patterns and quantify the similarities between an individual's repeated measurements. We can also estimate the learning rates (i.e., the slopes) of each group from GCM. We fitted the growth model in Bayesian framework by following the instruction by Oravecz and Muth (2018).

Analyzing repeated measurements poses a significant challenge, as accounting for the dependence within a person's data is crucial. Such dependence can violate statistical assumptions of data being independent and identically distributed. The GCM addresses this issue by adaptively modeling grouped data to account for longitudinal dependency. The model assumes that each measurement includes a noise component centered on the underlying growth curve, following a specific distribution. When the growth curve accurately depicts the genuine underlying trend, the estimations won't be tainted by systematic changes over time.

The GCM can be implemented within a hierarchical Bayesian framework, which allows the use of prior distributions to express our knowledge about the most likely values of model parameters. After specifying the priors, we fit the Bayesian model by calculating the posterior distribution of the parameters, representing the updated probability distribution conditioned on the data. Bayesian modeling's strength lies in its ability to integrate prior knowledge of parameters and derive the posterior distribution from the data without relying on the data collection plan.

In the GCM, an individual's unique growth trajectory is represented as a mathematical function illustrating the relationship between variables over time. We determine the growth function by examining the shape of the learning curve, as depicted in Fig.2-5. Based on the figure, we assume a linear growth curve model. $Y_{i,t}$ represents the accuracy of participant i at block t. In a simple growth curve model, we can express the within-person change over blocks in terms of intercept (initial level) and slope (rate of change) parameters. We can effectively model their learning trajectories by fitting a straight line to each participant's five blocks, with the x-axis representing time and the y-axis denoting accuracy rate. The GCM is specified as follows:

$$Y_{i,t} \sim N(\beta_{i,1} + \beta_{i,2} T_t, \sigma_e^2)$$

$$\beta_{i,1} \sim N(\mu_{\beta_1}, \sigma_{e_{\beta 1}}^2)$$

$$\beta_{i,2} \sim N(\mu_{\beta_2}, \sigma_{e_{\beta 2}}^2)$$

The first equation captures the effect of time at the person level and specifies the likelihood function. The distribution of $Y_{i,t}$ is Gaussian with the mean as a linear grown function and variance. In contrast, the second and third equations are group-level equations, which capture between-person variability in intercepts and slopes. Both parameters have the shape of the normal distribution, with their mean capturing the population mean intercept (or slope) and variance representing the individual differences.

To specify prior distributions for model parameters, we want to ensure that the prior is wide enough to fully cover the plausible range of the data. We specified the following priors, parametrized in terms of mean and variances:

$$\mu_{\beta_1} \sim N(0,1)$$

$$\mu_{\beta_2} \sim N(0,1)$$

$$\sigma_{e_{\beta_1}} \sim unif(0,1)$$

$$\sigma_{e_{\beta_2}} \sim unif(0,1)$$

We implemented the GCM within a Bayesian framework using R (with RStudio, RStudio Team, 2015) and JAGS, which are open-source statistical software packages. JAGS offers flexible and customizable extensions to accommodate various prior specifications. Table 2-4 presents the results of each condition obtained by fitting the GCM to the experimental data. By comparing the intercept values for water, heat, and car conditions in the Base task, we can observe that the initial values for the three groups are similar: water intercept (M=0.51, 95% HDI = (0.47,0.56)), heat intercept (M=0.53, 95% HDI = (0.46,0.6)), and car intercept (M=0.5, 95%

HDI = (0.43,0.6)). This suggests that participants' performance in the first block was not significantly different, as they began the tasks with comparable levels of information. Examining the contrasts between water-heat, water-car, and car-heat reveals that the ranges of all the differences include 0, indicating no significant differences among the three groups at the start of the Base task.

Base task	**Mean**	**95% HDI Low**	**95% HDI High**	**ESS**
water intercept	0.5128	0.4654	0.5601	6993
water slope	0.0419	0.0294	0.0540	5748
heat intercept	0.5277	0.4624	0.5956	6835
heat slope	0.0476	0.0299	0.0658	4383
car intercept	0.4978	0.4344	0.5612	2641
car slope	0.05	0.0323	0.0668	1870
water - heat (contrast) intercept	-0.0147	-0.0966	0.0658	7660
water – car (contrast) intercept	0.0152	-0.0643	0.0938	2806
car - heat (contrast) intercept	-0.0299	-0.1194	0.0642	3504

Table 2-4. The estimated intercepts and slopes of each condition from fitting GCM for the Base task

Crucially, when comparing the intercept values of the Target task presented in Table 2-6, the intercepts for the analogy conditions (i.e., water to heat, M= 0.71, 95% HDI = (0.66,0.78); heat to water, M=0.68, 95% HDI = (0.61,0.74)) are significantly higher than those for the non-analogy conditions (i.e., water control, M=0.51, 95% HDI = (0.47,0.58), heat control, M = 0.57, 95% HDI = (0.54,0.61)) in the target task. This result indicates that participants in the analogy conditions demonstrated better performance at the outset and that knowledge was successfully transferred from the Base task.

The estimated slopes represent the learning rate of participants in each group. As shown in Table 2-5, the estimated slopes in the Base task are similar, with values around 0.04, and their 95% HDI does not include 0. This indicates that participants' learning abilities were comparable, regardless of the information they received. In the Target task, the estimated slopes for all four conditions are similar but smaller than those in the Base task.

Target task	**Mean**	**95% HDI Low**	**95% HDI High**	**ESS**
water to heat intercept	0.7033	0.6372	0.7694	20540
water to heat slope	0.0206	0.0005	0.0407	19522
car-to-heat intercept	0.5353	0.4836	0.5879	18781
car-to-heat slope	0.0261	0.0018	0.0504	18898
heat to water intercept	0.6778	0.6028	0.7528	9120
heat to water slope	0.0260	0.0073	0.0333	8659
car-to-water intercept	0.5129	0.4372	0.5886	4642
car-to-water slope	0.0147	0.0001	0.0293	2296

Table 2-5. The estimated intercepts and slopes of each condition from fitting GCM for the Target task

The decreasing learning rates for both analogy and non-analogy groups can be attributed to different reasons. The analogy groups began with higher accuracy rates because they were able to use the corresponding cue validities derived from the mapping between the Base and

Target tasks. For instance, the accuracy rates for the water-to-heat condition increased from 0.70 in the 6th block to 0.82 in the 10th block. On the other hand, the non-analogy groups started with lower accuracy rates as no knowledge was transferred from the Base task. Consequently, their performance consistently remained lower than that of the analogy groups across all blocks, reaching around 0.67 in the last twenty trials.

In summary, the performance gap between the two groups demonstrates the knowledge transfer from the Base task to the Target task in the analogy group. Throughout the five blocks in the Target task, the analogy group consistently exhibited superior performance compared to the non-analogy group.

Conclusion

Previously, we emphasized the significance of understanding how individuals make decisions in new or unfamiliar environments. Our research centered on the role of analogical transfer in decision-making, shedding light on the potential benefits of employing past knowledge to navigate changing situations.

Our findings illustrated that participants' performance improved as they consistently engaged in probabilistic inferences within the same task environment. The results revealed that individuals tend to adopt diverse inference strategies and progressively adapt to the task environment by learning the most effective strategy, such as the take-the-best heuristic in this instance. Starting with slightly higher accuracy than random guessing, participants in both conditions exhibited increased accuracy rates over the five decision trial blocks. Furthermore, we identified the influence of analogical transfer in the analogy conditions. Participants who had previously learned the analog in the base task effectively mapped the corresponding cues to the target task. This mastery of cue relationships and cue validities conferred an advantage as participants commenced the target task equipped with more information and a higher accuracy rate. This advantage was maintained throughout all five blocks in the target task.

The study bridges the research on analogical transfer and heuristics, two fields that have not been previously interconnected. The experimental results compellingly demonstrate that analogical mapping could be a crucial process for organizing cues before applying heuristic strategies. This integration of research areas provides a more comprehensive understanding of decision-making processes and the role of past knowledge in enhancing decision-making accuracy.

In the following chapter, we will investigate the strategy selection processes throughout the trials by applying machine learning techniques. This approach will deepen our understanding of the role of analogical transfer in shaping strategy selection and adapting to decision-making contexts.

Chapter 3

Machine Learning Strategy Identification

We hypothesize that analogy can impact decision-makers' performance strategy selection. Specifically, drawing on past knowledge through analogies can improve decision accuracy by identifying similarities between past and current tasks. Additionally, by recognizing the connections, decision-makers are more likely to use the strategy that works best for the previous task. To investigate the impact of analogy on strategy selection, we developed a machine learning strategy identification method (Fang, Schooler, and Luan, 2022). This method allowed us to uncover the strategies used in the experiment and compare the performance of decision-makers who used analogies to those who did not. Compared to existing strategy identification methods (e.g., Glöckner 2009; Lee & Gluck, 2019), the machine learning approach proposed by Fang et al. (2022) provides several improvements. It incorporates a wider range of features that are extracted from participants' behavior, and uses a more advanced classification algorithm, resulting in higher accuracy.

In Chapter 2, we conducted Experiment I to investigate the impact of using analogies on decision-making tasks. Participants were randomly assigned to one of two conditions. In the analogy condition, they were first asked to choose which water flow system had a higher flow rate, followed by an analogous task of choosing which heat flow system had a higher heat flow rate. The two tasks were counterbalanced by swapping the order. In the no-analogy condition, participants were given an unrelated car efficiency task that did not conceptually relate to the subsequent heat flow task (or water flow task). Our results showed that participants in the analogy condition performed better overall than those in the non-analogy condition, with a higher accuracy rate at the beginning of the target task that persisted throughout all decision trials.

This chapter investigates how decision-makers change their strategies when adapting to a new task and how they select strategies when transitioning to an analogous task. We employ a machine learning approach to reveal the strategies used throughout the task and test the hypothesis that the most effective strategy from the previous task will be carried over to the subsequent task. The chapter is organized as follows[1]: First, we review commonly applied strategy identification methods in the literature. Next, we introduce the machine learning approach termed "Machine Learning Strategy Identification" (MLSI). In order to show the performance of MLSI, we compare it with the multiple-measure maximum likelihood (MM-ML) method. Finally, we discuss the performance of the MLSI approach within the context of our analogy experiment.

When conducting research on multi-attribute decision-making, researchers face the challenge of drawing inferences from behavioral data that shed light on cognitive strategies. This problem has been tackled by various methodological approaches, namely Structural Modeling (SM), Process Tracing (PT), and comparative model fitting.

One approach is to posit that people are equipped with a variety of strategies that they can select adaptively to solve the decision problems they face (e.g., Beach & Mitchell, 1978; Gigerenzer & Selten, 2002; Lieder & Griffith, 2017; Payne et al., 1993), and the selection of a particular strategy depends on many factors, such as information cost (Beach & Mitchell, 1978; Bröder, 2000), feedback on decision outcomes (Rieskamp & Otto, 2006), and cognitive cost

[1] The literature review and Experiment II of this chapter is updated from Fang et al., (2022)

associated with each strategy (Fechner et al., 2018). With a repertoire of possible strategies, it has been challenging for researchers to identify which strategies people use in a given task.

Some strategy identification methods use only individuals' choices to infer the strategy they may have applied (e.g., Bröder 2003; Bröder & Schiffer 2003; Hilbig & Moshagen, 2014; Lee, 2016). Other methods consider additional data, such as confidence ratings, decision time, protocols , and eye-tracking measures (e.g., Glöckner, 2009; Glöckner & Herbold, 2011; Rieskamp & Hoffrage, 2008; Lee et al., 2019). Compared to relying on choices alone, strategies can be identified more reliably and accurately when a variety of behavioral measures are taken into account (e.g., Glöckner 2009; Lee et al., 2019; Riedl et al., 2008).

The machine learning method for strategy identification (MLSI) considers choices and other behavioral data. Compared to existing methods, this method has fewer constraints on the form and the amount of behavioral data required, and it can identify strategies on a trial-by-trial basis. Thus, it can detect dynamic changes in strategy selection that elude other methods. Before introducing the method, we review the primary existing methods that have been employed to address the issue of identifying the strategies people use. We differentiate between outcome-based methods, which rely solely on the choices made, and process-based methods, which incorporate behavioral data leading up to the decision.

Outcome-based Methods

Choice outcomes are frequently used to infer decision-making processes because they are readily observed. For instance, in *structural modeling, researchers* run multiple regressions between cues and choice outcomes and take regression weights to indicate how heavily a person relies on each cue (e.g., Brehmer, 1994; Stewart, 1988). Because it only describes the statistical relations between cues and choice outcomes, this approach does not reveal much about the actual process of making a decision (Bröder 2000).

In *comparative model fitting*, a metric, such as maximum likelihood, is calculated to gauge how well a strategy describes choice outcomes. For example, Bröder and Schiffer (2003) compared an individual's choice outcomes with predictions made by several strategies and treated the strategy with the highest estimated likelihood as the one most likely to have produced the observed choice pattern. This method assumes that an individual uses only one strategy and applies that strategy with a constant error rate across different combinations of alternatives, which they referred to as item types. For this method to work well, researchers need to carefully design a set of item types so that the strategies make markedly different outcome predictions across trials (Jekel et al., 2011). The comparative model fitting method is not unique in this regard, diagnostic items are required for all identification methods to a greater or lesser extent. Hilbig and Moshagen (2014) developed this method further by using a multinomial process tree formalism, which allows varying error rates across item types (i.e., each type can have its own error rate) instead of a fixed error rate for all item types. The error rates are further grouped into random application errors and systematic errors associated with an item type, helping to increase the accuracy of strategy identification.

The comparative model fitting approach assumes that individuals use the same strategy over time and identifies strategies based on the overall choice outcomes in a task. However, studies have found that people may use a mixture of strategies over a sequence of decisions (e.g., Davis-Stober & Brown, 2011; Scheibenenne et al., 2013) and adapt strategies in response to environmental changes (Lee et al., 2014; Lieder & Griffith, 2017; Rieskamp & Otto, 2006). A latent mixture model can accommodate these findings because it allows for the possibility of incorporating multiple strategies into a single mixture model. Such models are amenable to various statistical methods, including Bayesian methods. For example, Scheibenenne et. al. (2013) used a Bayesian framework to infer what strategy or combination of strategies was most likely to have produced the outcome data by comparing the posterior probabilities of a single-strategy model and a multiple-strategy model (see also Lee, 2011, 2016).

Process-based Methods

Outcome-based methods draw inferences about strategies based on observed decisions. In contrast, process-based methods infer strategies from an array of dynamic process data associated with each strategy. The process data can be collected when *information is acquired, integrated, and evaluated. The methods to collect process data include mouse tracing, verbal reports, brain imaging, eye tracking, and so on.* Researchers analyze these process measures, inferring individuals' cognitive processes and thereby, the decision strategies used. Various process-tracing techniques have been developed to collect data about how people acquire information (see Schulte-Mecklenbeck et al., 2017 for a recent overview).

For example, information boards display cues and cue values in a matrix and are commonly used to track how individuals acquire information in an experiment (e.g., Bettman et al., 1990; Johnson, et al., 2008; Payne et al., 1993). In most computer-based information matrices (e.g., MouseLab), cue values are hidden behind boxes. At the beginning of a decision trial, all boxes in the matrix are closed. As a participant moves the mouse over or clicks on a box, it opens to reveal the cue value. Figure 3.1 shows an example of how participants may move their mouse on an information board display when applying take the best.

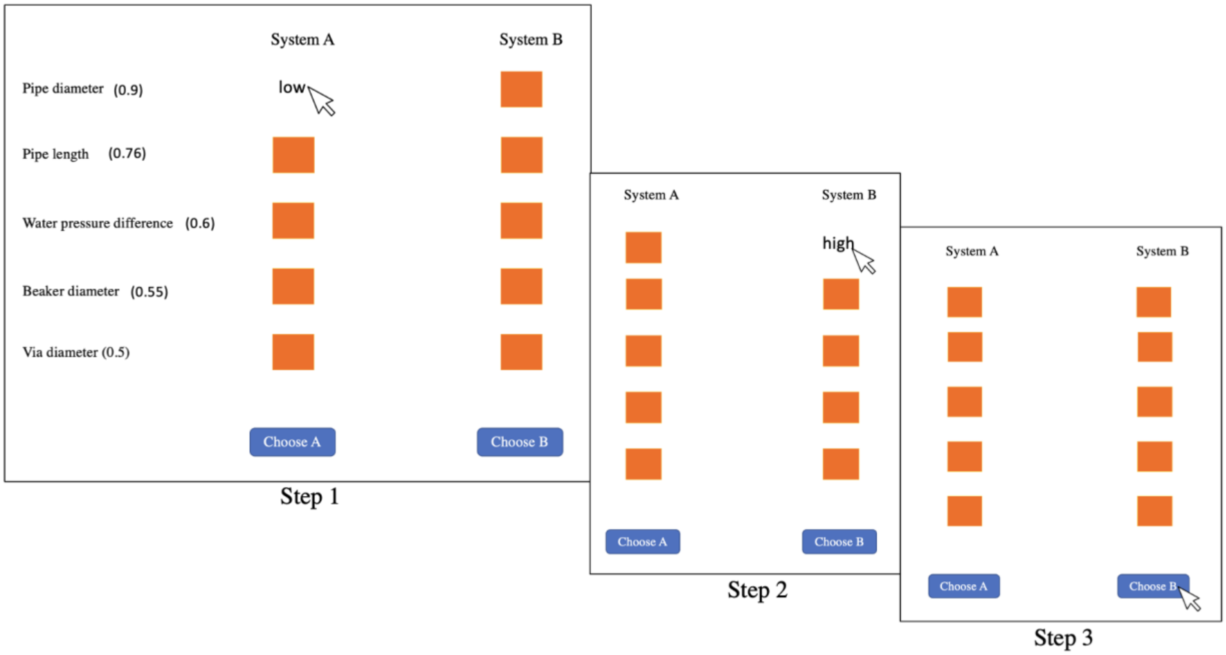

Figure 3-1 An example of how participants might move their mouse as they use take-the-best to make a decision based on cues displayed on an information board. In Step 1, a participant clicks on the box of the highest-ranked cue (i.e., the pipe diameter cue) for System A to reveal the cue value (i.e., low), and then in Step 2, clicks on the same cue for System B. Because values of the pipe diameter cue discriminate between the systems, the participant in Step 3 chooses System B, which is inferred to be the system having higher water flow rate.

A variety of process measures can be constructed from mouse movements on an information board that is indicative of different decision strategies. These measures include the total time spent on a trial, the proportion of information searched, variability in the amount of information searched per alternative, and the ratio of cue-wise transitions to alternative-wise transitions, to name a few (for a comprehensive list, see Riedl et al., 2009). These measures have also been extended to eye-tracking studies (Krol & Krol, 2017; Schulte-Mecklenbeck et al., 2017).

The conventional approach to strategy identification compares the observed process measures to the canonical process patterns of that strategy (e.g., Day, 2010; Glöckner & Herbold, 2011). Sometimes, analyses of process measures are used to bolster conclusions from outcome-based methods. For instance, Rieskamp and Hoffrage (2008) found that participants who were identified as using take-the-best or WADD based on their choices differed significantly on six process measures. When different strategies make the same choices, process measures provide more information to assist strategy identification.

Combining Outcome-based and Process-based Methods

Strategies make predictions about not only the decisions people make, but also the information search and cognitive processing they undertake before reaching decisions. Researchers have stressed the importance of combining outcome-based and process-tracing data to uncover human decision processes more accurately than using each type of data alone (Costa Gomes et al., 2001; Harte & Koele, 2001; Glöckner 2009; Lee et al., 2019). For example, Riedl et al. (2008) developed a decision tree, named DecisionTracer, with three process measures and one outcome-based measure as the decision nodes. Glöckner (2009, see also Jekel et al., 2010) developed the multiple-measure maximum likelihood (MM-ML) method that integrates outcomes, decision times, or confidence ratings to identify strategies based on the Bayesian information criterion.

Incorporating multiple sources of data may decrease the number of decisions required to identify the strategies people are using. Several recent attempts have been made to identify strategy switches because decision-makers can use different strategies over time, even when facing similar decisions (Lee & Gluck, 2020). Brusovansky et al. (2018) proposed a model that deploys strategies in a trial-by-trial stochastic manner by using a probabilistic switching parameter. Lee et al. (2019) incorporated decision outcomes, verbal report data, and search behavior into a Bayesian hierarchical model to infer when individuals may have changed strategies and how often they have done so. The existing methods make inferences about strategy switch based on observations of multiple trials. However, in many situations outside the laboratory, people do not make repeated decisions. Here, we use machine learning techniques to identify strategies on a trial-by-trial basis by integrating process and outcome data collected in one decision trial.

Machine Learning Strategy Identification (MLSI)

The problem of strategy identification entails inferring individuals' strategies based on behavioral data, such as choice outcomes and process measures. These data help differentiate strategies because each strategy is presumably associated with a signature data pattern. Machine learning (ML) techniques, specifically supervised learning, solve similar classification problems. The goal of supervised learning applied to strategy identification is to find a good classifier that can distinguish between strategies based on the combination of outcome and process data in data sets where strategy labels are known[2]. The resulting classifier is then applied to assign strategy labels in data sets where individuals' strategies are unknown. We next explain the detailed workflow of this method (see Figure 3-2 for an overview).

[2] We use identification and classification interchangeably throughout the chapter.

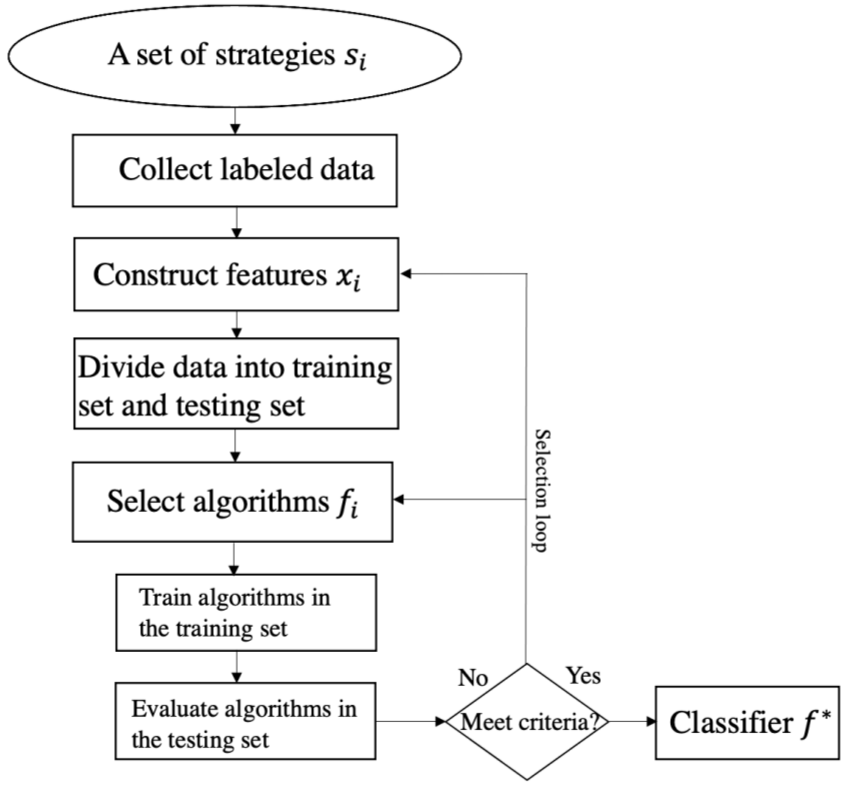

Figure 3-2 The process of identifying strategies using machine-learning algorithms.

1. *The Strategy Identification Problem*

The ML method aims to find a function that maps behavioral traces (i.e., all data collected during an experiment) to strategies. We look for this function by training an ML algorithm on a data set with known trace-label associations. The trained model will then be able to assign a given behavioral trace with a strategy label, such as take-the-best or WADD.

2. *Collect Labeled Data*

The labeled data are relevant behavioral data engendered by a strategy, and their forms depend on data collection (e.g., mouse movements or eye tracking). To increase the efficacy of trained ML algorithms, the labeled data should represent the unlabeled data that will need to be classified later.

3. Construct Features

Even a single decision can produce a large amount of raw data. However, using the raw data to differentiate strategies is not ideal because these data typically contain much noise and irrelevant information. Moreover, the high dimensionality of the raw data can result in computational inefficiency. Therefore, it is crucial to derive a set of informative features from the raw data when building a classification model. For example, the process-based approaches we reviewed above suggest potentially useful features, such as the time spent reading information and the proportion of information searched. Furthermore, using meaningful features can help improve the interpretability of the resulting classifier and, in turn, provide a better understanding of the strategies people use to make decisions. With these goals in mind, we constructed features from the raw data.

4. Divide Data into Training and Testing Sets

Labeled data are divided into a training set and a testing set. An ML model is trained on the training set, while its performance is evaluated on the testing set. A model's performance in the testing set is one way to measure its ability to generalize to unseen data, and one standard metric to evaluate model performance is its identification accuracy on the testing set.

Some ML models have hyperparameters that control the learning process (e.g., the maximum depth of trees in Random Forest). These hyperparameters need to be tuned to optimize a model's performance. A common approach is testing all hyperparameter combinations in a predefined search space through K-fold cross-validation. Specifically, the training set is split evenly into K subsets. A model has trained on $K-1$ subsets and evaluated on the remaining subset, and the procedure is repeated for each subset to find the best combination of hyperparameters.

5. Select ML Algorithms

To classify decision strategies, we use classic supervised learning algorithms, including K Nearest Neighbors (KNN), Random Forest (RF), Decision Trees (DF), Support Vector Machines (SVM), and Multilayer Perceptron (MLP). The algorithms are implemented in Python using the scikit-learn library (Pedregosa *et al.*, 2011). A comprehensive review of these and other ML algorithms can be found elsewhere (e.g., Friedman et al., 2001; Kotsiantis, 2007).

The performance of an ML model is determined mainly by how accurately it classifies the testing set. If the performance does not meet preset criteria, we will return to the previous stage, trying to improve the diagnosticity of the features or include more ML algorithms (see Figure 3-2). One criterion is a model's relative performance compared to other models. We also consider the interpretability of a model and the ease of collecting the required behavioral data.

A Worked Example

Here, we present an example of how the ML approach works for strategy identification, using KNN and an SVM model with a linear kernel (Linear SVM) as the classification models (see Figure 3.3). We first simulated 40 take-the-best and 40 WADD trials. In each trial, we labeled the trial with the strategy that generated its data and recorded two features, total decision time and proportion of information searched. The 80 trials were divided so that 60% of the trials were used for training and 40% for testing. The ML models learned the decision boundaries that separated the labeled trials based on the training set. A trial from the testing set was labeled according to which side of the boundary it was located. The learned decision boundary can be linear or nonlinear, depending on the ML algorithm. In our case, Linear SVM builds a linear decision boundary, whereas KNN produces a nonlinear boundary.

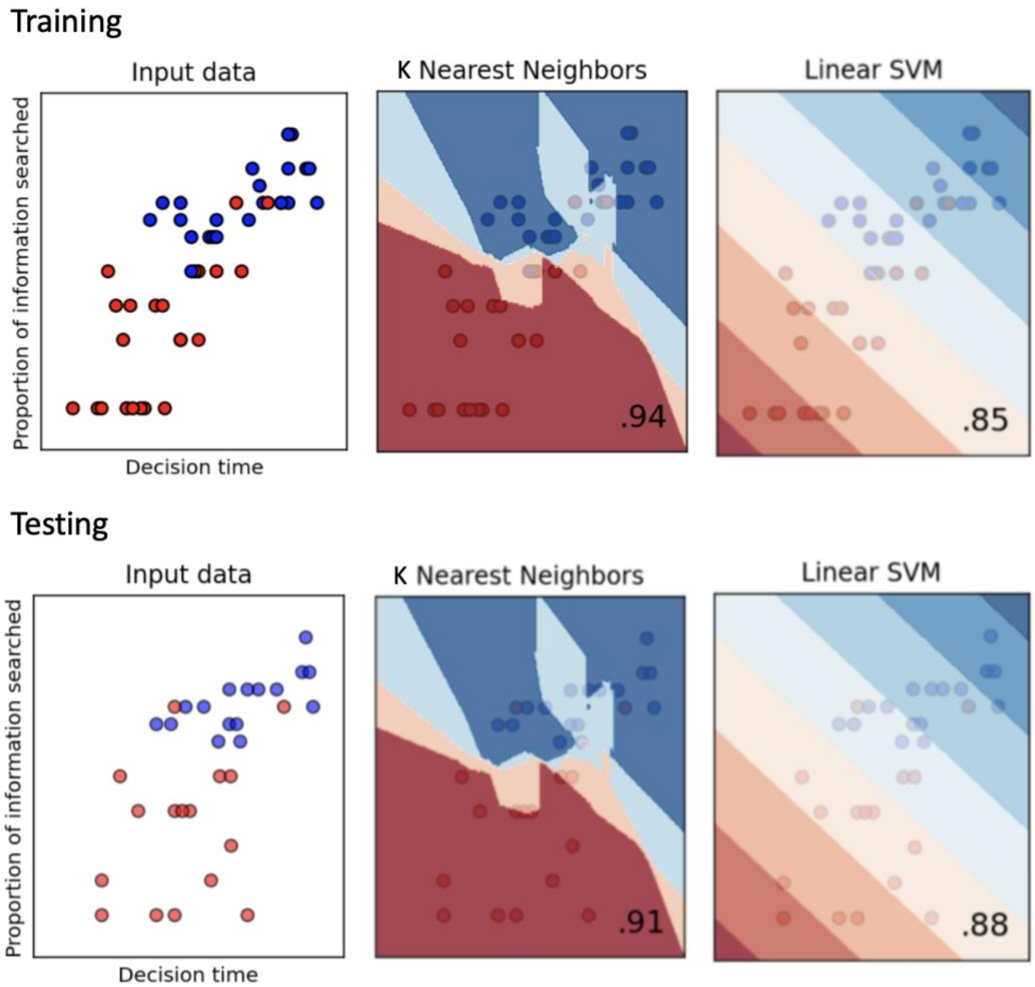

Figure 3-3 Simulated data for take-the-best (red dots) and WADD (weighted-additive; blue dots) and the accuracy of two machine-learning algorithms. Each dot represents data from one decision trial. In the two panels in the leftmost column, the *x*-axis is the decision time in seconds, and the *y*-axis is the proportion of information searched; these were the two features processed by the machine-learning algorithms. To test the accuracy of the algorithms, 60% of all data were used for training and 40% for testing. The leftmost column shows data from the training set (top) and the testing set (bottom). The upper panels in the two right columns show the classification boundaries constructed by K Nearest Neighbors and Linear SVM, respectively, based on the training set data, and each model's training accuracy is shown in the lower right corner. The lower panels in these columns show model predictions in the testing set data, and the identification accuracy is shown in the lower-right corner. The color shadings indicate confidence in model identifications, with darker shadings representing higher confidence levels.

In conclusion, we have comprehensively reviewed strategy identification techniques that employ choice outcomes, process measures, or a combination of both. We have also unveiled a novel machine learning strategy identification (MLSI) approach, demonstrated with a practical example. Following this, we apply the MLSI method to decipher strategies people adopt in our experiment. A comparative analysis between MLSI and the MM-ML method is conducted to evaluate the effectiveness of our newly introduced MLSI approach. Finally, we assess the transferability of these strategies under the analogy condition.

Strategies Identification Experiments

We evaluated the MLSI method by investigating how well it identifies individuals' decision-making strategies in multi-attribute tasks. In the first experiment, we pitted the MLSI method against the MM-ML (Multiple-Measure Maximum Likelihood) method, comparing their respective strategy identification performance. Our findings demonstrated parity between the two methods, underscoring the reliability of MLSI. However, MLSI distinguished itself through its unique ability to disclose strategies with a heightened degree of precision, allowing for a trial-by-trial analysis. This granular insight offers a more nuanced understanding of the evolution of decision-making strategies during tasks. In the subsequent experiment, we applied the MLSI method to the analogy experiment, as outlined in Chapter 2. This primarily aimed to examine our hypothesis, asserting that analogy significantly influences strategy selection. Through these investigations, we aimed to not only assess the efficacy of MLSI but also elucidate the intricate dynamics governing multi-attribute decision-making strategies.

Experiment II

In this experiment, we aim to validate the efficiency of MLSI by comparing it with MM-ML in terms of strategy identification accuracy. We instructed participants to use specific decision-making strategies, namely take-the-best, Weighted Additive Rule (WADD), or Tallying, after which they proceeded to make a sequence of decisions based on their assigned strategy. We employed the stimuli from Glöckner (2009)'s study to guarantee the generation of distinctive predictions on the dependent variables by each of the strategies, which is a prerequisite for the effective application of MM-ML. It is important to note that we excluded Δ-inference from this experiment. This was because all the cue values in this task environment were binary, and in such a scenario, the performance of the take-the-best strategy and Δ-inference is virtually indistinguishable.

Participants

A total of sixty undergraduate students were recruited from the participant pool of the Psychology department at Syracuse University. These students were then randomly assigned to one of three predefined strategy conditions, with each condition having an equal distribution of 20 participants. Prior to the commencement of the experiment, informed consent was duly acquired from each participant. The duration of the experimental session was approximately half an hour. In return for their participation, students were rewarded with research credits, which contributed towards the fulfillment of their course research requirements. It's important to note that three participants from the WADD condition were unable to complete the experiment. After excluding these individuals from the study, the resulting dataset included data from 57 participants.

Materials

Participants were asked to take the role of a college student interested in purchasing a used car with the aim of choosing the car that would last longer. Each cue took either a favorable (1) or an unfavorable (0) value, with favorable values associated with more durability. Each car was described by four cues: mileage, model year, the number of previous owners, and the number of accidents. The cues were ranked by their validities, which were shown to the participants. MM-ML requires that different strategies make different predictions on multiple dependent variables. To meet this requirement, we used stimuli from Glöckner (2009, Table 1) containing six stimuli types. The six types were repeated ten times each, resulting in 60 trials for each participant. The order of decision trials was randomized for each participant.

Procedure

Participants were trained to apply a certain strategy and made decisions using a computerized information board. Every participant went through a tutorial on how to use one of the strategies and was given five practice trials. Feedback was provided for each practice trial. Participants needed to make correct decisions on at least 80% of the practice trials to proceed to the decision task; otherwise, they repeated the tutorial from the beginning. After successfully learning the strategy, participants engaged in a 60-trial decision task in which their decision outcomes and mouse movements were recorded. A "correct" decision was defined as a choice in agreement with the strategy they were trained on. Participants were given one point for each correct decision, and their goal was to maximize total points. There was no monetary reward given to the participants.

Results

The descriptive statistics shown in Table 3-1 were in line with our expectations for the strategies. A Bayesian one-way ANOVA analysis indicates that the total decision time and the proportion of information searched were significantly different among the strategies (BFs > 1,000).

Group	*N*	Rate of strategy adherence	Decision time (in s)		Proportion of cues searched	
			M	SD	M	SD
TTB	20	0.946	3.81	0.35	0.37	0.02
WADD	17	0.830	9.74	2.02	0.98	0.02
Tallying	20	0.964	6.02	0.66	0.98	0.01

Table 3-1 Descriptive Statistics for Experiment II

Construct features

The feature selection step aims to generate features from the labeled raw data, enabling ML algorithms to classify trials as originating from users employing take-the-best, tallying, or WADD strategies. we selected the 21 features shown in Table 3.2. The first two features (x_1, x_2) are the decision time and the proportion of cues searched. Because there were four cues in this experiment, we reduced the numbers of features for the time to read each box (x_3 ... x_{10}) and those to code search order (x_{11} ... x_{18}). Three features (x_{19} ... x_{21}) record the final choice outcomes for each strategy.

Feature number	Definition
x_1	Decision time
x_2	Proportion of cues searched
$x_3 \dots x_{10}$	Time needed to process cue values shown in Boxes 1 to 8
$x_{11} \dots x_{18}$	Eight features to record search orders
$x_{19} \dots x_{21}$	Whether the choice outcome is line with take-the-best, weighed-additive (WADD), or Tallying.

Table 3-2 Feature Set for Machine-Learning Models in Experiment II

MLSI Analysis

We first analyzed the performance of MLSI at the trial-by-trial level. We then compared the performance of MLSI and MM-ML at the individual level.

Strategy identification at the trial-by-trial level

We randomly selected five participants from each strategy group to form the testing data set. For the remaining 42 participants, we used 10-fold cross-validation to train the ML models and search for optimal hyperparameters. Each participant made 60 decisions; therefore, there were 2,520 training trials and 900 testing trials. We applied five ML models in the training set and evaluated their performance in the testing set. Table 3-3 shows the identification accuracy of these ML models in the testing set. MLP performed the best, with an identification accuracy of 91.8%.

ML models	Random forests	SVM	KNN	Decision tree	MLP
Experiment II	91.3%	91.4%	91.2%	87.4%	91.8%

Table 3-3 Strategy Identification Accuracy of Machine-learning Models on Test Participants

We report the results of Random Forest in more detail, because Random Forest is easier to interpret than SVM, in terms of feature importance, and its identification accuracy of 91.3% is very close to that of MLP. We plot the trial-by-trial identification results by Random Forest in Figure 3-4. It shows that Random Forest was best at discriminating take-the-best from Tallying and WADD, yielding perfect identification accuracy for take-the-best. The identification accuracy for WADD and Tallying was 85.6% and 88.3%, respectively. Because the search patterns of WADD and Tallying are similar, the majority of the misclassifications were between WADD and Tallying. The identification accuracy also differed among participants. For example, the trained Random Forest model mistakenly classified 26 of Participant 9's trials and perfectly classified six participants' trials (Participants 8, 11, 12, 13, 14 and 15).

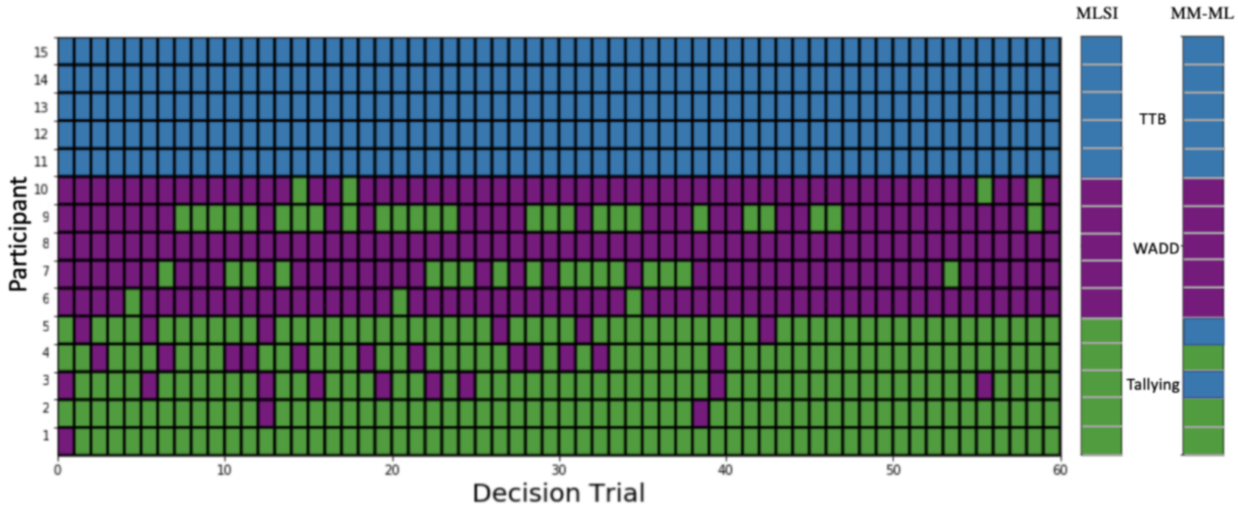

Figure 3-4 Results of strategy identification for test participants in Experiment II. Strategies were identified by Random Forest, the best-performing machine-learning algorithm. The five participants at the top were trained to use take-the-best (TTB), the middle five weighted-additive (WADD), and the bottom five Tallying. The colored boxes indicate strategies predicted by Random Forest for each participant in each decision trial, with blue for TTB, purple for WADD, and green for Tallying. The two columns on the right show the overall strategy identified by the machine-learning strategy identification (MLSI) method and the multiple-measure maximum likelihood (MM-ML) method, respectively, for each participant.

Feature importance

In addition to having an accurate model, it is beneficial to have an interpretable one as well. In strategy identification, we also want to determine which features are crucial for distinguishing strategies. A deeper understanding of the model's logic can help confirm its accuracy and potentially enhance the model by selecting the most relevant features. While the random forest algorithm can be challenging to interpret due to its randomized nature, it is still possible to glean insights about the most important features of the random forest.

Gini importance represents the average decrease in impurity that each feature contributes across all trees in the forest (Archer & Kimes, 2008). By calculating the Gini importance for each feature, we gain insight into which features have the most significant impact on the model's predictions. This information can be valuable for understanding the underlying logic of the model and identifying the most relevant features for strategy identification. Additionally, it can help simplify the model by removing less important features and possibly improving its performance by focusing on the most influential ones.

Figure 3-5 shows the Gini importance of the 21 features. The time needed to read the bottom boxes of the first alternative (Box 4, x_6) and the second alternative (Box 8, x_{10}), and the total decision time (x_1) were important in differentiating take-the-best, Tallying, and WADD. A likely explanation is that WADD participants generally spent more time at the bottom boxes, because they would need to take some time to integrate cue values and calculate the overall score of an alternative. The search order features (x_{11} ... x_{18}) were also important, because they indicate whether a participant used cue-wise or alternative-wise search.

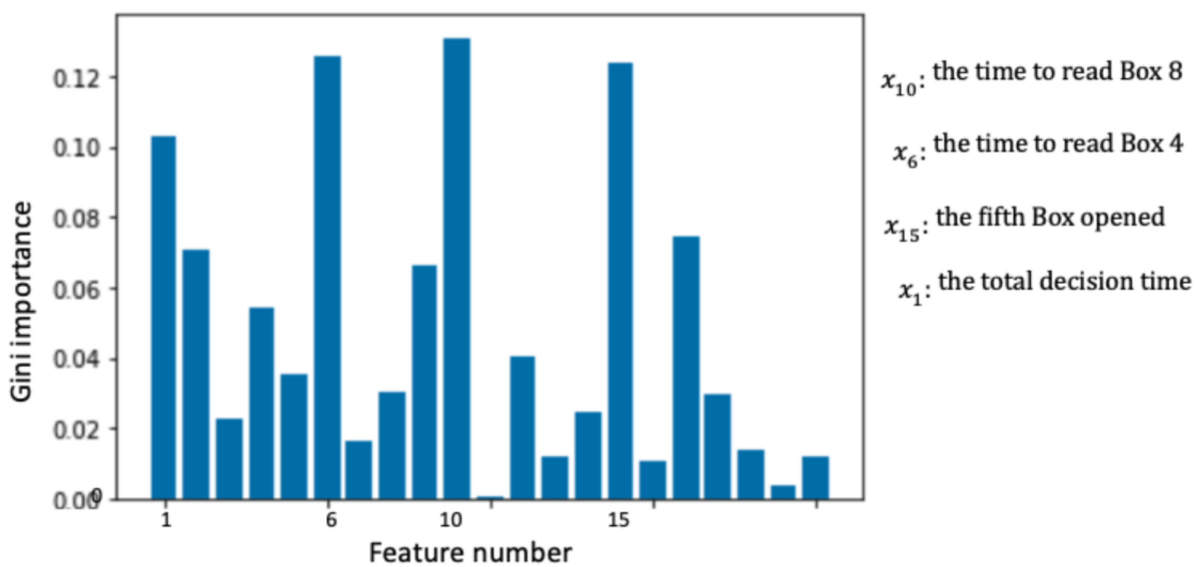

Figure 3-5 The Gini importance of each feature in Experiment II. Gini importance measures how much Random Forest relies on a particular feature in strategy identification. The sum of Gini importance is 1.

Classification by participants

We aggregated the trial-by-trial classification results for each participant to classify the participant as being best described by one of the three strategies. This was done for MLSI based on how the majority of the trials for a participant were classified. Figure 3-4 shows that Random Forest perfectly classified the fifteen test participants at the individual level.

MM-ML Classification

We used the R code provided by Jekel and colleagues (2010) to conduct the MM-ML analysis. Using a combination of decision outcomes, decision times, and confidence judgements, MM-ML estimates the likelihood that a participant used a particular strategy. We ran MM-ML with participants' choice outcomes and decision times. With these two types of process data as inputs, MM-ML classified 86.7% of the participants correctly. Figure 3-5 shows that of the 15 test participants, two Tallying participants (i.e., Participants 3 and 5) were misclassified as using take-the-best. MM-ML never misclassified a take-the-best participant as using WADD, and vice versa. MM-ML's classification performance might have been even better had we collected confidence ratings.

Discussion

We have shown that MLSI can identify the strategies participants used for each trial with high levels of accuracy. The most frequently misclassified trials were between WADD and Tallying, because they are both compensatory strategies and result in similar search patterns. The identification accuracy varied among the participants. For example, the identification accuracy was relatively low for Participant 9. A plausible reason is that we trained a model based on data of a group of participants and applied that model to classify the idiosyncratic behavior of new participants. A future study could investigate if the identification accuracy of a participant would improve by using a model trained on data of the same participant. At the individual level, MLSI perfectly classified each test participant, and MM-ML accurately classified 86.7% of the test participants. That said, each approach has its own strengths and weaknesses. MM-ML classifies participants at the individual level using carefully designed stimuli but does not need training data, whereas MLSI can classify at both the trial and the individual levels on a broad range of stimuli but does require training data.

We also applied MLSI to an existing dataset (Walsh and Gluck, 2016) and compared the classification results to Lee & Gluck (2019), which used Bayesian methods to identify when participants changed strategies on the same dataset. Both methods strongly agree with predicting participants using TTB, but there is also disagreement for participants with mixed strategies. Please see Appendix A for details.

Experiment III

In the analogy experiment described in Chapter 2, participants in each condition undergo three tasks. The first task, referred to as the base task, involves participants learning the cues and the best-performing strategy in one of three conditions (i.e., water flow, heat flow, and car efficiency). Participants then proceed to the second task based on their assigned conditions. In the analogy condition, those who learned in the water flow task are assigned to the heat flow task, and vice versa. In the car control condition, participants who learned in the car efficiency task were randomly assigned to either the water flow or heat flow task. The third training task involves training all participants in a specific strategy. The number of training trials is determined by their performance in the base and target tasks.

In this experiment, we aim to apply MLSI to the main tasks by using the labeled data from the training task. The goal is to recover the strategies used throughout the experiment and, specifically, we are interested in the transfer of the strategies at the beginning of the target task for the analogy condition.

Collect labeled data

Labeled data were collected during the training task in the experiment. The total number of training trials is 180; however, if participants make a correct decision in the base and target tasks, the number of training trials is reduced by one. To demonstrate how the reward system operates, we display Figure 2-4 in the instructions provided to the participants. The more accurate decisions a participant makes in the base and target tasks, the fewer training trials they will receive. Participants are rewarded for spending less time in the experiment by performing well in the first two tasks.

The training strategy is randomly assigned to each participant, regardless of the best-performing strategy in the task experiments. Table 3-4 presents the total number of training trials for each strategy. The total number of trials varies because participants' performance differs, resulting in individual variations in the training trials.

Heuristic strategies	**Take the best**	**Δ-inference**	**Tally**
# of labeled training trials	1812	1534	1687

Table 3-4 The number of training trials for each strategy.

Construct features

The feature selection step aims to generate features from the labeled raw data, enabling ML algorithms to classify trials as originating from users employing take-the-best, tallying, or Δ-Inference strategies. Participants interacted with the experiment interface, depicted in Figure 3-6. Their search behaviors were recorded as mouse coordinates every five milliseconds. Figure 3-6 displays the mouse traces of three representative trials where take-the-best, Δ-Inference, and tallying were applied to the same decision pair. These traces demonstrate the cue-wise search characteristic of Δ-Inference and take-the-best and the alternative-wise search expected from tallying. Subsequently, we encoded the mouse trace data as features indicative of the respective strategies.

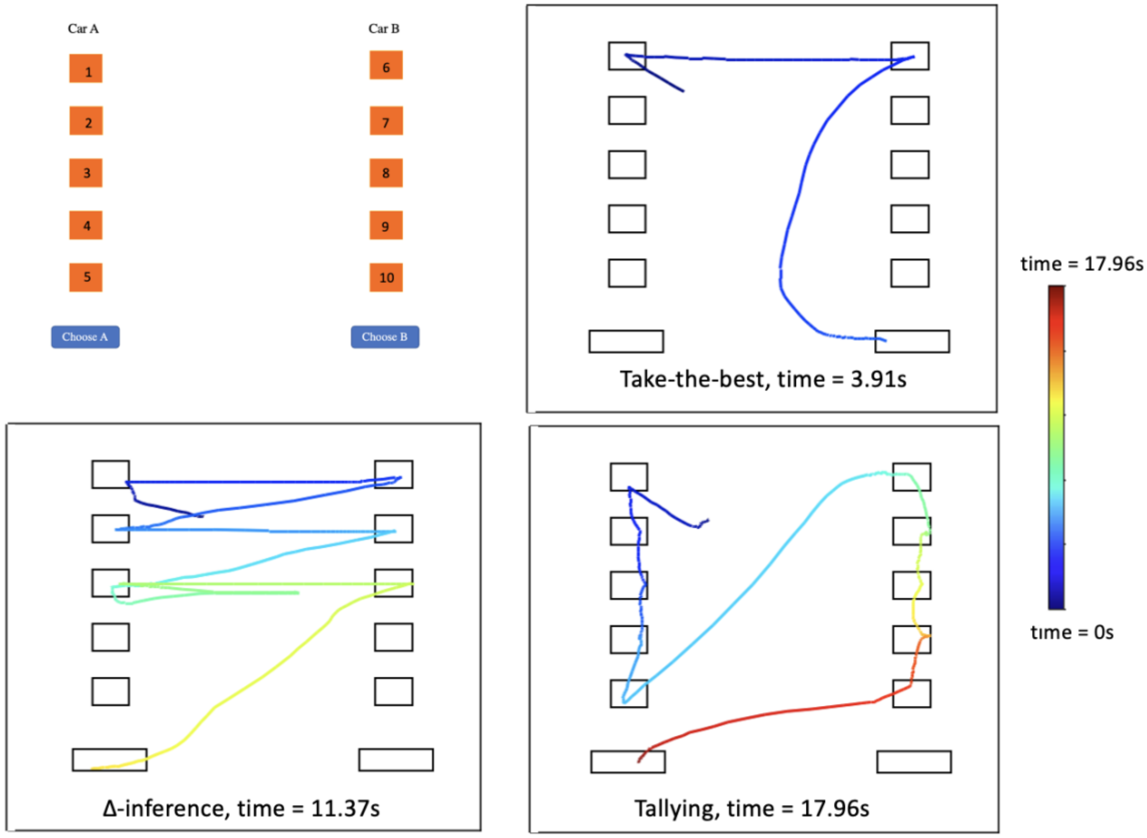

Figure 3-6 Examples of mouse movement paths for participants trained to use take-the-best, Δ-inference and tallying respectively. In the experiment, participants must click on boxes to see cue values. The numbers on the boxes in the upper left panel are used to identify features based on the mouse movement paths (see Table 3-5), which were not shown to the participants. The colors of a movement trace indicate how much time had elapsed since the trial started. The total time participants spent on the trial is shown next to the labels of the trained strategies.

Table 3-5 summarizes the features that were selected. A primary feature is the decision time in each trial (x_1). As discussed previously, noncompensatory strategies tend to search for less information, resulting in shorter decision times. Compensatory strategies integrate all cue values, which potentially takes longer. The proportion of cues searched (x_2) reflects the number of boxes that have been opened. Participants using noncompensatory strategies typically need to open fewer boxes than those using compensatory strategies.

Feature number	Definition
x_1	Decision time
x_2	Proportion of cues searched
$x_3 \dots x_{12}$	Time needed to process cue values shown in Boxes 1 to 10
$x_{13} \dots x_{22}$	Ten features to record search orders
x_{23}	Whether the choice outcome is in line with take-the-best
x_{24}	Whether the choice outcome is in line with Δ-inference
x_{25}	Whether the choice outcome is in line with Tallying
x_{26}	Whether alternative values are the same in the cue before the discriminating cue

Table 3-5 Feature Set for Machine-learning Models in Experiment III

Each box on the information board was assigned a number. We recorded the total time taken to process each box, yielding ten features (i.e., x_3 to x_{12}) for a five-cue decision task. Ten features ($x_{13} \dots x_{22}$) represent the search order; that is, the order in which the boxes were opened. We recorded search order by entering the box number of each opened box into one of these ten features. If fewer than ten boxes were opened, then zeros were recorded in the remaining search order features. Noncompensatory and compensatory strategies are associated with cue-wise and alternative-wise search, respectively. An example of take-the-best's search order is (1,5,2,6,0,0,0,0,0,0) and Tallying is (1,2,3,4,5,6,7,8,9,10).

Participants' choices were compared with the predictions of each strategy (x_{23} to x_{25}). If a participant's decision was consistent with a particular strategy, the feature value was coded as a "1," whereas an inconsistent decision was coded as "0." For example, if a participant's final choice is consistent with take-the-best and Δ-inference but not with Tallying, the feature vector is (1,1,0). An additional feature x_{26} is designed to differentiate take the best from Δ-inference. It is a binary feature that indicates whether the cue values on the cue before the discriminating cue were the same. If they were different, the participant was likely using Δ-inference, and the feature is encoded as "0"; otherwise, the participant was more likely using take-the-best, and the feature is encoded as "1." The intuition behind this feature is that when the cue values for both alternatives differ from the previous cue, the participant could not have used take-the-best because they would have already decided on this cue.

MLSI Analysis

We first analyzed the performance of MLSI at the trial-by-trial level using the labeled training dataset. We used 10-fold cross-validation to train the ML models and search for optimal hyperparameters. We applied five ML models in the training set and evaluated their performance in the testing set. Table 3-6 shows the identification accuracy of these ML models in the testing set. MLP performed the best, with an identification accuracy of 92.2%.

ML models	Random forests	SVM	KNN	Decision tree	MLP
Experiment III	90.67%	86.3%	88.16%	83.83%	92.2%

Table 3-6. The MLSI results in 10-fold cross-validation.

In this case, we use the best-performing trained model, MLP, to uncover strategies used in the base and target tasks of the two analogy groups.

Strategies in the base and target tasks

Since we do not have access to the ground truth of participants' strategies in the base and target tasks, as they were not trained to use any of the presumed strategies, these strategies remain unlabeled. We employed the trained model to assign labels to each trial. The first step in this process involves extracting the same set of features used for training ML models from the base and target tasks. We calculated 26 features for each trial in both tasks, as outlined in Table 3.3, and then used the trained model to identify the strategies employed. Figure 3.7 presents the classification results for each block in the analogy conditions. Each bar represents the proportion of trials per participant in the block using the corresponding strategy.

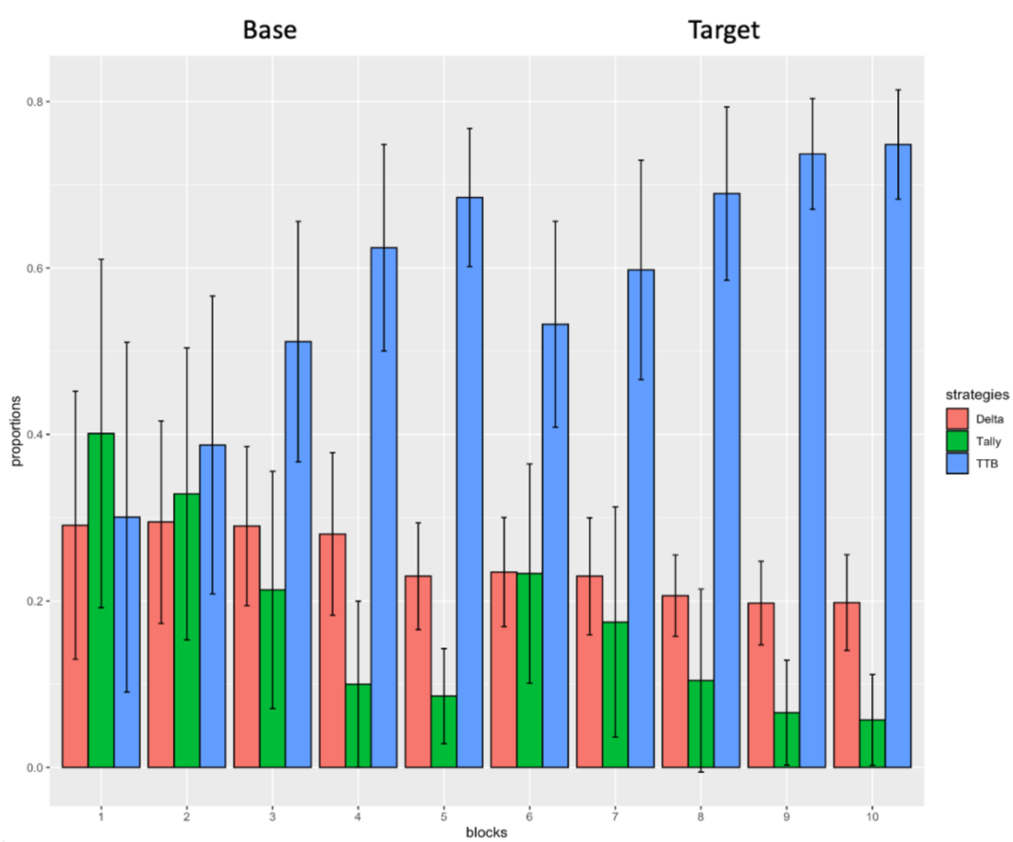

Figure 3-7 The classification results for each block. Each bar indicates the proportion of trials in the block using the corresponding strategy. Each strategy is color coded as indicated in the legend.

Two trends can be observed from the results. First, the proportion of Take-the-Best (TTB) usage increases for each block, as the choice outcomes of TTB are consistent with the feedback provided to participants. In other words, TTB is the best-performing strategy in the task environments. Second, the proportion of compensatory strategies decreases because they do not offer participants accurate feedback and require considerable time and effort to apply.

Feature Importance

Figure 3-8 shows the Gini importance of the 26 features (Table 3-5). The feature designed to differentiate between Δ-inference and take-the-best (x26) is the most important one. The second most important feature (x14) codes the box opened second. If participants applied cue-wise search, the second box opened should correspond to the highest validity cue of the second alternative (box 6). In contrast, an alternative-wise search suggests that the second box opened should correspond to the second highest validity cue of the first alternative (box 2). Moreover, as suggested by the analysis of decision times, total decision time (x1) is also an important feature. The time to read the second alternative's highest validity cue (x8) could also help distinguish Δ-inference from take-the-best because it generally takes longer to assess whether two cue values differ by a Δ than to determine whether they are different.

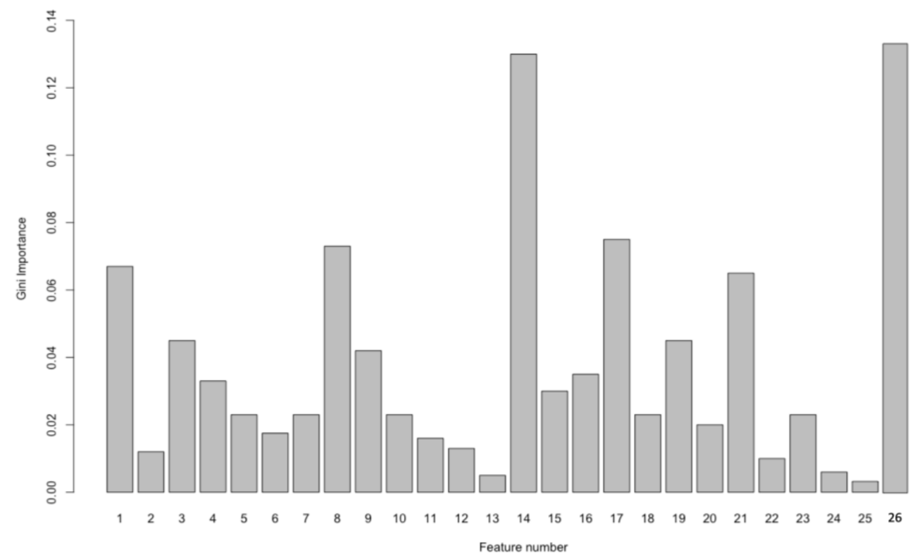

Fig. 3-8 The Gini importance of each feature. Gini importance measures how much Random Forest relies on a particular feature in strategy identification. The sum of Gini importance is 1.

Conclusion

In this study, participants received feedback based on the decisions that Take-the-Best (TTB) would have made. We analyzed the data from this experiment using an MLP model trained on data from the training phase. The Machine Learning Strategy Identification (MLSI) analysis demonstrates that participants adapted to the task environment according to the feedback they received. This result is also consistent with the findings of Rieskamp and Otto (2006). Participants initially tried a variety of strategies, but most ultimately converged to the non-compensatory strategies TTB and Δ-inference, which often led to similar decisions in this environment.

We can identify strategies on a trial-by-trial basis, which is an advantage over some strategy identification approaches that require the strong assumption that decision-makers employ a single strategy throughout a sequence of decision tasks (e.g., Bröder, 2003). With a rich set of features reflecting the characteristics of different strategies on a single trial and ML algorithms' ability to learn the relationship between these features and the strategies, we can identify strategies in a single decision trial based on a decision-maker's search behavior and choice.

The ability to identify strategies on a trial-by-trial basis is expected to assist researchers in investigating factors that influence strategy selection, particularly in this study focusing on the impact of analogical transfer. We have demonstrated knowledge transfer from the perspective of strategy selection. The majority of strategies identified as using TTB is indicative of the knowledge transfer from the base task. This finding suggests that the transferred knowledge consists of the cues' importance and how they are utilized, revealing valuable insights into the decision-making process and the role of analogical transfer in shaping strategy selection.

Chapter 4

Modeling Strategy Selection and Analogical Transfer in ACT-R

We demonstrated that individuals acquire strategies through feedback and subsequently apply these strategies to succeeding tasks by recognizing analogous components. To delve deeper into the fundamental cognitive mechanisms in play, we aim to enhance our exploration of this phenomenon by creating a computational model grounded in the Adaptive Control of Thought-Rational (ACT-R) framework (Anderson et al., 2004) in this chapter. Our choice to use the ACT-R as our preferred modeling framework is motivated by two key reasons. First, our modeling objective hinges on two primary components: strategy selection and analogical transfer. ACT-R is a cognitive architecture capable of integrating models of these two facets of decision making, thereby offering a comprehensive portrayal of the pertinent cognitive processes (Anderson et al., 2004). Second, prior modeling endeavors in ACT-R have focused on strategy selection and analogical mapping, providing a solid base to merge both components into a singular cognitive model, thereby shedding light on the intricate cognitive dynamics that enable strategy learning and analogical transfer.

Using ACT–R to Model Decision Strategies

ACT-R is an extensive, quantitative cognitive theory that elucidates a broad array of cognitive phenomena. These phenomena encompass areas such as memory performance (Anderson et al., 1998; Borst & Anderson, 2013), strategy selection (Nellen, 2003; Marewski & Link, 2014; Dimov et al., 2017), multi-tasking abilities (Salvucci et al., 2008), cognitive skill acquisition and transfer (Singley & Anderson 1989; Taatgen 2013; Anderson et al., 2021), and math problem solving (Anderson et al., 2011; Lee et al., 2016). Given the complexity of ACT-R theory, the introduction focuses on aspects pertinent to the current task -- strategy learning and analogical mapping. Please refer to Anderson (2007) for a comprehensive overview.

Modelers lean on at least two fundamental mechanisms, termed modules, to construct a cognitive task model within the ACT-R framework. These modules correspond to the two types of memory: declarative memory and procedural memory. The declarative module pertains to declarative memory, storing factual knowledge, while the procedural module represents procedural memory, implementing production rules for various skills. In ACT-R, the fundamental unit of declarative memory is designated as the "chunk." Each chunk is an instantiation of a specific chunk type, and chunk type defines the kind of information that chunk can hold. Specifically, the chunk type lays out "slots" that can contain links to other chunks, thus setting up the structure of information within a chunk. Chunks can signify simple memories, such as the names of distinct objects (for instance, beaker, vial, pipe), or they can encapsulate more intricate knowledge regarding specific attributes of an object (such as the diameter of a pipe). For instance, to model the knowledge "pipe diameter located at line one on the monitor is a cue and has a validity of 0.9", we can declare a chunk type "cue" with slots "name", "location", and "validity". We then can create chunks and model the knowledge as follows:

Base-task-Pipe-diameter

isa	cue
name	pipe-diameter
location	line-one
validity	0.9

The "isa" refers to the chunk type of which the chunk "Knowledge-pipe" is an instantiation. The three chunk slots (name, location, validity) contain links to other chunks. For example, “line-one” is also a chunk that could describe the x-y coordinates of that location. The other cues, such as beaker diameter, vial diameter, and pipe length, are stored as cue chunks in the declarative memory. Each chunk is associated with an activation, which controls how likely and quickly a chunk can be retrieved. The activation values are affected by how often the chunk has been retrieved and how long ago these retrievals occurred.

Later, at the beginning of the target task. The ACT-R agent encodes the knowledge of cues of the target task with no cue validities from the screen. For instance, the cue “bar diameter” can be encoded as the following:

Target-task-Bar-diameter

isa	cue
name	bar-diameter
location	line-three
validity	?

The cue validity is not provided for the target task. Instead, it is the process of analogical mapping, transferring the cue validity from the base task and filling this information gap.

In contrast, procedural knowledge is modeled using production rules, also colloquially referred to as if-then rules. Each rule comprises conditions (the "if" component of the rule) that are evaluated against various chunks, such as the contents in declarative memory. When the conditions of a production rule are met, the rule is eligible for selection. If subsequently chosen, the actions in the "then" segment of the rule are executed. As productions fire, they produce a sequence of actions that can be compared to human actions in the same task.[3]

In the current study, wherein participants use a cursor to reveal information about options by clicking on boxes, and subsequently receive feedback about their choices, several additional modules become integral to the process, as depicted in Figure 4-1. These include the visual module, which controls visual attention and encodes the attributes (e.g., locations, color, value, etc.) of the visual stimuli on the monitor; the motor module, which controls the execution of physical actions and movements (e.g., move mouse, click button); and the goal module, which maintains objectives and anticipated outcomes. Further, the imaginal module manages internal representations and simulations of information not immediately perceptually accessible. Lastly, the declarative module comes into play, archiving pertinent information as chunks. These modules underscore the ACT-R's diverse functionality and comprehensive cognitive modeling potential, contributing to a holistic understanding of cognitive tasks such as those in this study.

---

[3] Existing production-system architectures such as ACT-R (Anderson & Lebiere, 1998), Soar (Newell, 1994), EPIC (Kieras & Meyer, 1997), and 3CAPS (Just & Carpenter, 1992) have been able to model human behavior in numerous tasks.

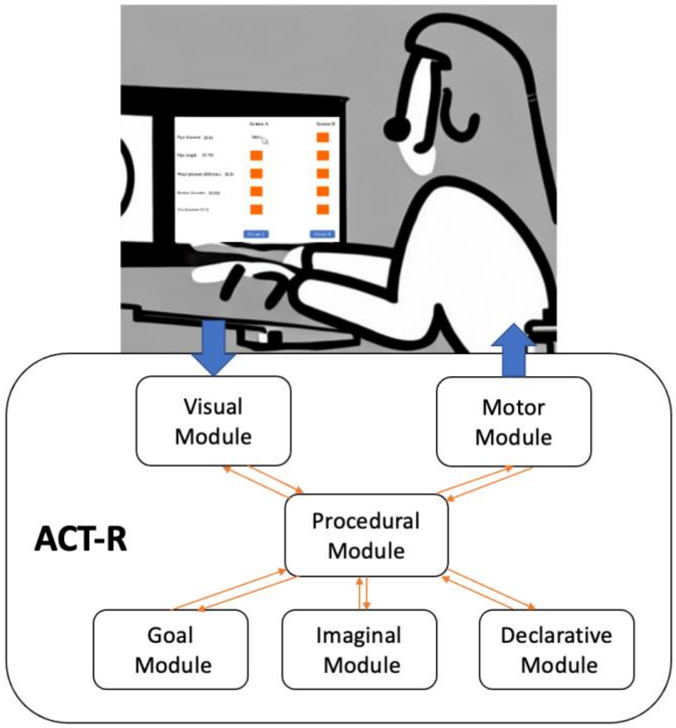

Figure 4-1. The Adaptive Control of Thought-Rational (ACT-R) cognitive framework is exhibited, highlighting the six ACT-R modules pertinent to the current study. Central to the modeling of decision strategies is the procedural module, which manages the production rules within procedural memory that orchestrates the activities of the other modules. The goal module maintains the representation of the decision goals that interact with the experimental stimuli. Meanwhile, the imaginal and declarative modules are employed to manipulate and preserve information relevant to the task. The visual module is engaged to perceive information necessary for the decision-making task, and the motor module is harnessed to select alternatives. This description encapsulates the interplay and integration of different modules within the ACT-R cognitive architecture, illustrating how they collectively contribute to the decision-making process.

The three heuristic strategies, take-the-best, Δ-inference, and Tallying, are operationalized as sequences of production rules within procedural memory (e.g., Fechner et al., 2016; Fechner et al., 2018; Marewski & Schooler, 2011; Schooler & Hertwig, 2005). These production rules interact dynamically with ACT-R's modules (as illustrated in Figure 4-1). Most of these modules have buffers that can store information in chunks. These chunks encode various forms of information, and they can be modified and used by the other modules. To illustrate, consider a participant who is moving the cursor to open a box covering an alternative's cue on a monitor, aiming to perceive and remember the cue's value. This entire process can be modeled within the ACT-R framework. Initially, the ACT-R agent sets a goal to navigate the mouse to the box's location, which involves the motor module. Following this, the box is opened, prompting the ACT-R agent to employ the visual module to identify the cue value and store it within a dedicated buffer. Concurrently, the imaginal module accesses the cue value stored in the visual buffer, encoding the information into a chunk. Finally, this chunk is stored in the declarative memory. The collaboration and interplay between different modules and buffers in ACT-R showcase how complex cognitive tasks, such as decision strategies, can be modeled and understood through this framework.

Decision Strategy Models

We have constructed ACT-R models that embody the general processing steps for take-the-best, Δ-inference, and tallying strategies. These models were specifically designed to undertake the same tasks as the participants in the empirical study, as detailed in Chapter 2. In the following section, we elucidate the processing steps of the model in relation to the tasks deployed in the study. A detailed flowchart of the take-the-best model is shown in Figure 4-2. Please refer to Appendix B for an in-depth exploration of the other models. Simulating the decision task in this study requires setting parameters that govern the modules' performance. Instead of tailoring the model parameters to fit the data, we predominantly used default or standard parameter values, as identified in previous research. This method has the potential to promote the generalizability and replicability of our models.

Take-the-best model

The procedural flow of the take-the-best model is illustrated in Figure 4-2. In each trial of the decision task, this model sets its goal to select which of the two water/heat flow systems exhibits a higher water/heat flow rate. Initiating with chunks, the model aims to capture data about cue names, locations, and validities. Furthermore, it maintains a chunk of task-specific information within the imaginal buffer, serving as the problem state. This particular chunk is deployed to compare cue values for a single attribute. The model set its goal to pinpoint the location of the most important cue – the cue with the highest validity – and directs visual attention towards it (boxes 1-3). To identify the most important cue, the model is programmed to randomly attend to a cue and register its validity and location. Once the model encodes all cue locations and validities, it initiates a comparison procedure to determine the cue with the highest validity and subsequently shifts its attention to that location (box 4). Upon locating the attribute, the model encodes the cue values of the alternatives. It sets its goal to find the attribute value location of the left alternative, redirects visual attention there, and stores the attribute value in the imaginal buffer (boxes 5-7). The model then repeats the same procedure for the right alternative (boxes 8-10). In a situation where the attribute values of the alternatives are identical, the model redirects its attention to the second most important cue (box 11) and replicates the actions carried out for the most important cue. Conversely, if the values on the currently examined attribute differ, the model opts for the alternative with the higher value, simulating a motor response akin to a mouse click (box 12). If the model exhausts all attributes without discovering one that distinguishes between the alternatives, it resorts to guessing (box 13).

Δ-inference model

The Δ-Inference model parallels the performance of the take-the-best model, with the primary distinction lying in the decision rule. The Δ-Inference model quantifies the disparities between two cue values and halts its search process when the difference exceeds 1 unit, such as in the instance of a high versus low comparison. When the difference between the cue values is precisely 1 unit, as seen in a low versus medium contrast, the model persists in exploring a lower-ranked cue. Please see Appendix B-1 for the detailed flowchart.

Tallying model

In each trial of the decision task, this model sets its goal to select which of the two water/heat flow systems exhibits a higher water/heat flow rate. Initiating with chunks, the model aims to capture data about cue names, locations, and validities. Furthermore, it maintains a chunk of task-specific information within the imaginal buffer, serving as the problem state. This specific chunk is used to store all cue values and aggregate these values for both alternatives. The model assumes a value of 1 to "high" and treats all other cue values as 0. Consequently, once all cue values have been encoded, the model initiates the process of summing these values for the alternatives. Upon completing this process for the first alternative, the model performs the same operation for the second alternative. Lastly, the model favors the alternative with the higher score. In cases where both alternatives yield the same final score, the model resorts to guessing. Please see Appendix B-2 for the detailed flowchart.

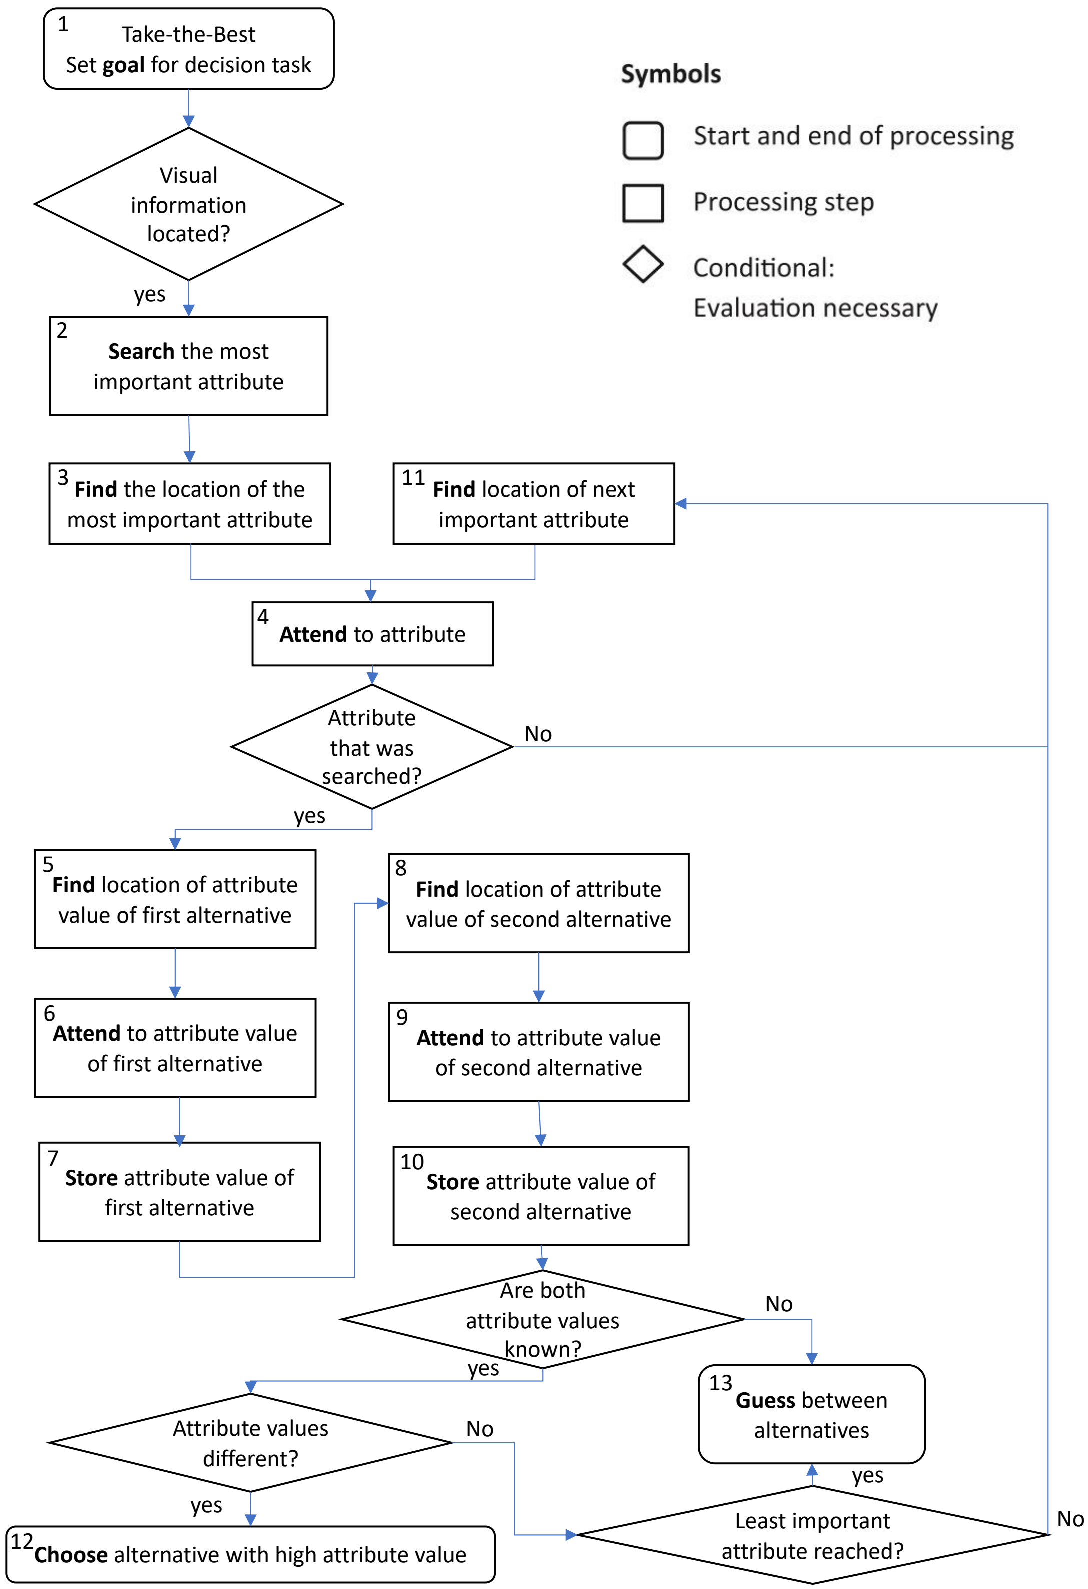

Figure 4-2. Flowchart of the ACT-R model for take-the-best

Learning from Feedback

After implementing the three strategic models, the subsequent step involves configuring the learning mechanisms that respond to feedback and foster the most effective strategy for the task. The cost-benefit paradigm is widely used to elucidate the strategy selection process in decision-making (Beach & Mitchell, 1978; Payne et al., 1988, 1993; Rieskamp et al., 2006). This paradigm interprets the strategy's costs as the time and effort expended for its selection, while the benefits encompass positive feedback or monetary gains. Decision-makers optimize these trade-offs by choosing the most efficient strategy from their repertoire, thereby striking a balance between accuracy and effort.

The role of learning is heavily emphasized in numerous quantitative theories of strategy selection (Busemeyer & Myung, 1992; Erev & Roth, 2001; Rieskamp et al., 2006). For instance, Rieskamp and Otto's (2006) model posits that individuals select from a range of strategies to gauge their accuracy. Similarly, the ACT-R cognitive architecture integrates reinforcement mechanisms that encourage the selection of cognitive processes (i.e., production rules) that underlie successful strategies (Fu & Anderson, 2006; Lovett & Anderson, 1996). Each production rule carries an associated utility that can be updated. Given a set of competing productions with expected utility values U_j, the probability of choosing production i is expressed by the formula:

$$Probability(i) = \frac{e^{U_i/\sqrt{2}s}}{\sum_j e^{U_j/\sqrt{2}s}}$$

In this formula, the summation j encompasses all the available productions for selection. The production with the highest utility will be the one chosen for execution. Utilities have noise, s, added to them. The noise is distributed according to a logistic distribution with a mean of 0. In the current study, we followed the ACT-R instruction and set noise s as 3.

As the model operates, it learns the utilities of productions based on the rewards it receives. These utilities undergo updates following a straightforward integrator model (Bush & Mosteller, 1955). Suppose $U_i(n-1)$ is the utility of production i after its $(n-1)$th application, and $R_i(n)$ denotes the reward the production i receives for its n_{th} application. In that case, its utility $U_i(n)$ after its n_{th} application can be expressed as:

$$U_i(n) = U_i(n-1) + \alpha[R_i(n) - U_i(n-1)]$$

Here, α signifies the learning rate, typically assigned a value of 0.2. This formula essentially embodies the Rescorla-Wagner learning rule (Rescorla & Wagner, 1972). The reward $R_i(n)$ that production i will receive will be the external reward received minus the time from production i's selection to the reward. This serves to give less reward to more distant productions. This reinforcement propagates to all the productions which have been selected between the current reward and the previous reward. As per this rule, the utility of a production undergoes gradual adjustments until it aligns with the average reward that the production receives.

Analogical Transfer via Path Mapping

Previously, we have implemented the reinforcement learning of three heuristic strategies within the ACT-R's built-in cost-benefit tradeoff mechanism, which is essentially the model replication of Rieskamp and Otto (2006) in ACT-R. We are now poised to integrate a version of analogical mapping to model the knowledge transfer process in the target task. One existing model, the path-mapping theory (Salvucci and Anderson, 2001) within ACT-R, implements this process effectively. This method allows us to identify and exploit relationships between concepts across two related domains. Choosing the path-mapping theory is not arbitrary but rather a matter of convenience for our proof-of-concept model. As a model of analogy already embedded within the ACT-R architecture, it offers an efficient and accessible mechanism for demonstrating our concept. Many theories of analogy have been successful in explaining numerous facets of analogical mapping, but they haven't provided a robust solution to the challenge of integrating such behavior with generic problem-solving tasks (Foburs et al., 2017).

Salvucci and Anderson (2001) seek to overcome this challenge by offering a systematic approach to building models of analogical problem-solving. Instead of progressively developing a broad cognition theory from an analogy theory, this theory pivots around the concept of mapping within the framework of a pre-existing comprehensive cognition theory (Anderson & Thompson, 1989; Salvucci & Anderson, 1998). This theory suggests that the key to such integration lies in how individuals amalgamate their knowledge and skills related to analogical reasoning with their expertise that is specific to the content domain under consideration. This approach promotes a cohesive fusion of analogical mapping with diverse problem-solving skills, thereby fostering a more holistic cognitive process.

The knowledge representations and how the analogs are mapped between knowledge representations are vitally important in the path-mapping theory. The representation component outlines how analog relations are manifested as different types of chunks in declarative memory, a process that takes place in the declarative module. The Path mapping segment delineates how the theory retrieves analogous paths of connected chunks to establish analogical mappings between higher-order structures, a process that hinges on the procedure module. Lastly, the model directs how models of analogical tasks accommodate task-specific processes. In this case, it is the transfer cue validities.

Representation

The Path-Mapping Theory represents analogs as higher-order structures, akin to other theories of analogy discussed in Chapter 1 (e.g., Gentner, 1983). This method achieves hierarchical structures by leveraging three key components in the knowledge representation (ACT-R's chunks): objects, relations, and roles. Objects serve as the semantic elements within the analogs, while relations function to interconnect objects or other relations based on their distinct roles. Roles are vital in forming a more complex knowledge structure by linking objects and relations. Figure 4-3 illustrates the knowledge representation of water flow and heat flow systems used in the ACT-R model, with objects and relations depicted as ovals and roles as rectangles. The chunks encode the information that "the beaker's pressure (object) is greater (relation) than the vial's pressure (object), causing (relation) water flow in the pipe (object)." The roles enclosed in rectangles not only connect objects to relations (as illustrated in the lower part of Figure 4-3) but also link lower-order relations to higher-order relations (e.g., 'greater' and 'cause'). Each role has five components: a parent slot, pointing to the parent relation; a parent-type slot, indicating the semantic type of the parent relation; a slot, indicating the relation slot

Role-Beaker-Attribute-diameter

isa	role
parent	beaker
parent-type	base-beaker
slot	attribute
child	diameter
child-type	base-diameter

that the object fills in the relation; a child slot, pointing to the child object or relation; and a child type slot, indicating the semantic type of the child object or relation. One example of the role "beaker, attribute, diameter" is in the lower left of Figure 4-3.

This decomposition process is critical in mapping relations from the base task to the target task.

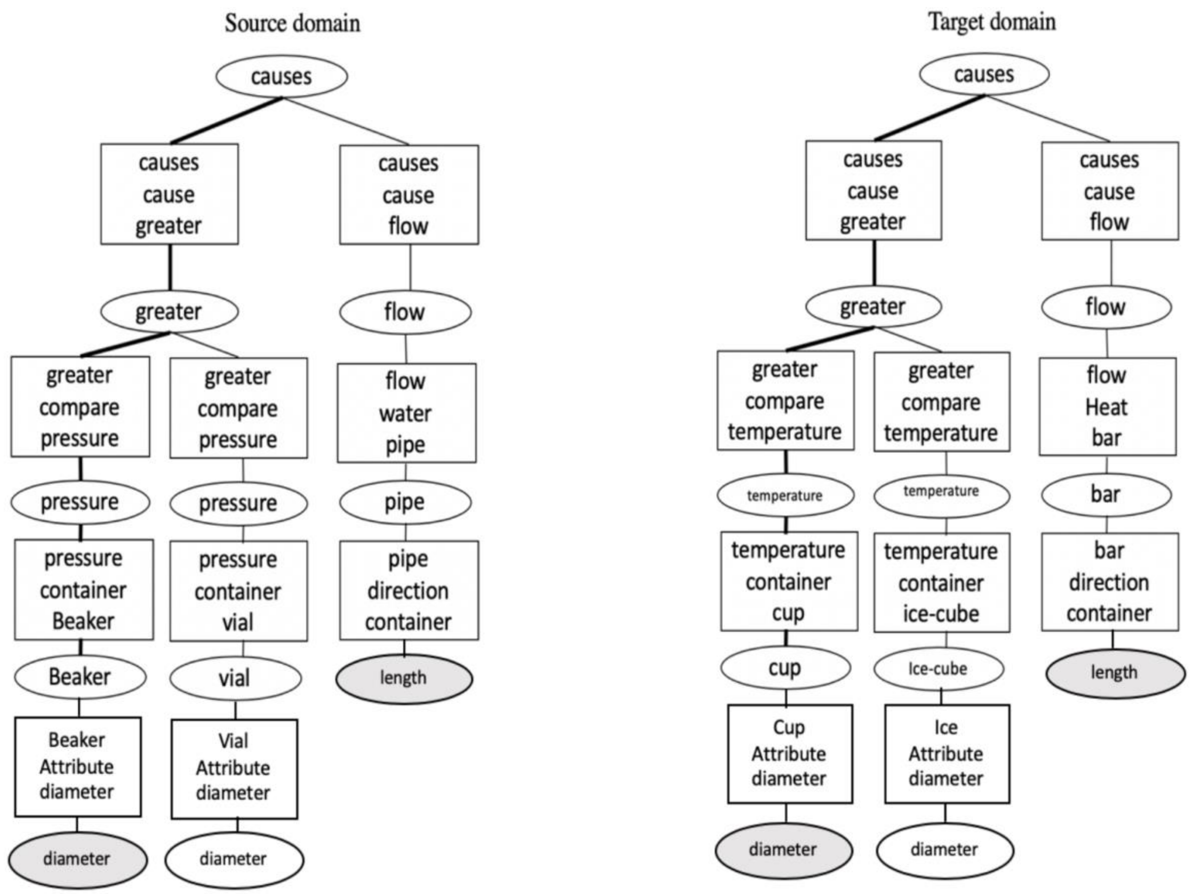

Figure 4-3 The ACT-R knowledge representations for water flow and heat flow.

Path mapping

Upon establishing the knowledge representations of the base and target tasks, a mechanism to find the analogous components between the two is required. This is accomplished through the Path-Mapping Theory, which aligns an object in the base task with a corresponding object in the target task by identifying similar paths between these objects and their highest-order relations. The process begins by defining a base task path from an object to its highest-order relation. It then maps the relations and objects along this path to analogous ones in the target task.

This mapping process starts from the highest-order relations, in alignment with the systematicity principle of the Structure-Mapping Engine (SME) that we discussed in Chapter 1 (Gentner et al., 1986). Using the example of water flow in Figure 4-3, the model begins with the object "diameter," shaded in grey. It retrieves its parent role "Beaker, Attribute, Diameter," and the model continues this iterative process of retrieving higher-level objects or relations until it reaches a point where no further relation can be obtained—essentially when it reaches the top relation. In this instance, "cause" is identified as the top relation. The path that extends from the bottom object to the top relation exemplifies the pathway identified by the path-mapping theory.

The subsequent step involves mapping each relation and object along the identified path in the base task to its corresponding analog in the target task. Commencing with the highest-order relation (in this instance, "cause"), path mapping seeks out the most similar target relation and establishes a map to it. This process proceeds further down the path in the base task, mapping each relation along the path to its respective analog, culminating in the final mapping between the lowermost objects in both base and target tasks. Take, for instance, an effort to map the "diameter" (highlighted in grey in Figure 4-3) to the target task. We would map the path identified in the initial step to an analogous path in the target task. This process initiates by retrieving the target relation that is most similar to the highest-order relation in the base task. In this scenario, path mapping would fetch "cause" in the target task, which is most similar to "cause" in the base task, thereby forming the mappings between the two tasks. This process continues recursively until the bottom objects are reached.

Simulation Results

Our study replicated the decision tasks from the empirical study that the participants completed using our simulation models. These tasks comprised two segments: the base task, involving 100 trials where participants were provided with cue names and their associated validities; and the target task, which comprised another set of 100 trials, but with no cue validities. In structuring the simulations, we made an assumption regarding participant behavior in the non-analogy conditions: we posited that participants ranked the cues based on their physical locations, viewing the topmost cue as the most important and the bottommost as the least. This supposition helped streamline the modeling process, given the substantial variability observed in participant responses under non-analogy conditions. In the absence of cue validities, their approach to using cues could differ widely: some participants might choose cues based on their sequence from top to bottom, while others might form their own analogical mappings. Given this complexity, simulating individual behavior with no assumptions in non-analogy conditions would be impractical. See Supplementary materials for the simulation results from ACT-R agents.

In the analogy condition for the target task, the transfer of knowledge extends beyond the learned cue validities; it also includes preferences for strategies. The analogical mapping process facilitates the transference of cue validities to the target task. As for the strategy preferences, understood as the utilities linked with each strategy, we initialized these at the start of the target task to match the learned utilities observed at the conclusion of the base task. This approach effectively allows for the transference of strategy knowledge to the target task.

Figure 4-4 displays the simulation results, specifically the accuracy rates of using the take-the-best, Δ-inference, and tallying models. These results illustrate the precision with which each model selected the alternative predicted by its strategy without accounting for any execution errors.

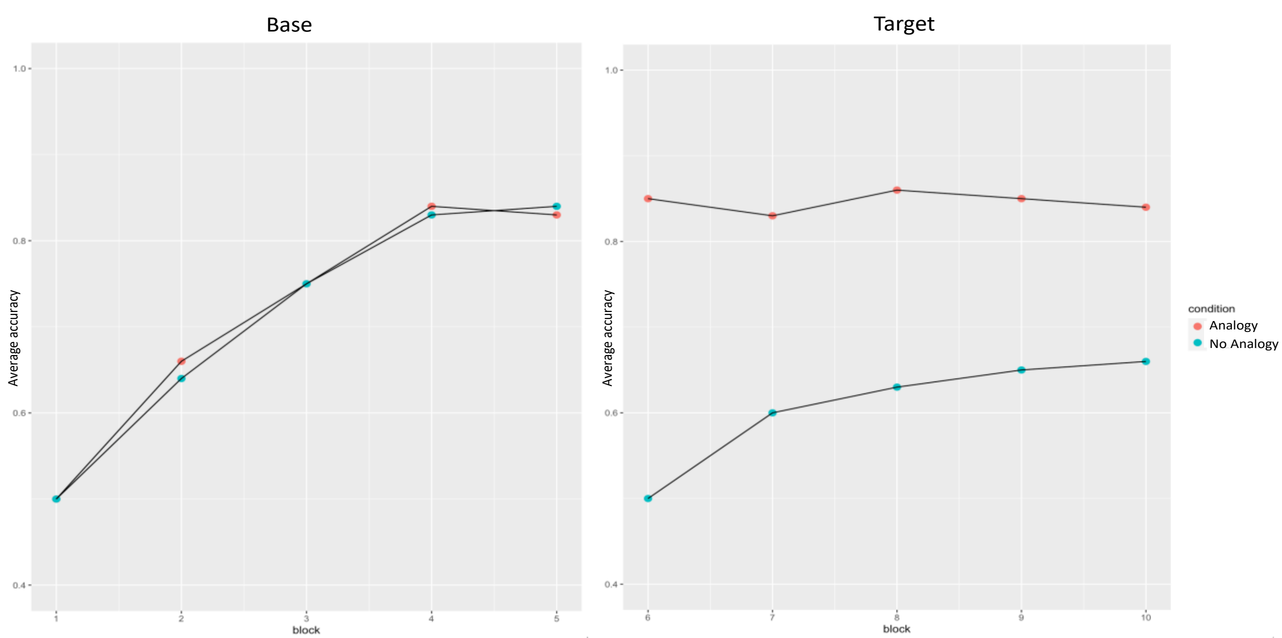

Figure 4-4 Results from the AC-R simulations when assuming participants used take-the-best, tallying, and Δ-inferences. Each block records the decision accuracy. The simulated ACT-R agents update their strategy utilities every time they receive feedback.

Figure 4-4 demonstrates an increase in decision accuracy across successive blocks within the base task, stabilizing during the final two blocks. This pattern is similar to the behavioral data seen in Figure 2-5, albeit converging towards a more consistent level of accuracy and reaching the ceiling more quickly. In the behavioral experiment, participants exhibited a continuous learning of strategies, with their learning curve persistently rising across ten blocks. Conversely, the simulated learning curve peaked by the fourth block, after which the performance plateaued and remained stable. The noted divergence could be traced back to two key factors. Firstly, our model did not account for potential errors during the strategy application phase. Secondly, we employed a default learning rate of 0.3 throughout the strategy learning period. By simulating the experiment under these ideal conditions, the ACT-R agents are expected to learn at an accelerated pace.

The simulation results also present another interesting observation: a slight decline in decision performance upon transition to the target task, as noted in the human subjects' data in Figure 2-5. This downward trend, however, was not reflected in the simulation. To account for this observed performance dip when transitioning to a new task, we would need to incorporate an additional component within the current model. This highlights potential avenues for refining our model to capture the dynamics of decision-making performance across different tasks more accurately.

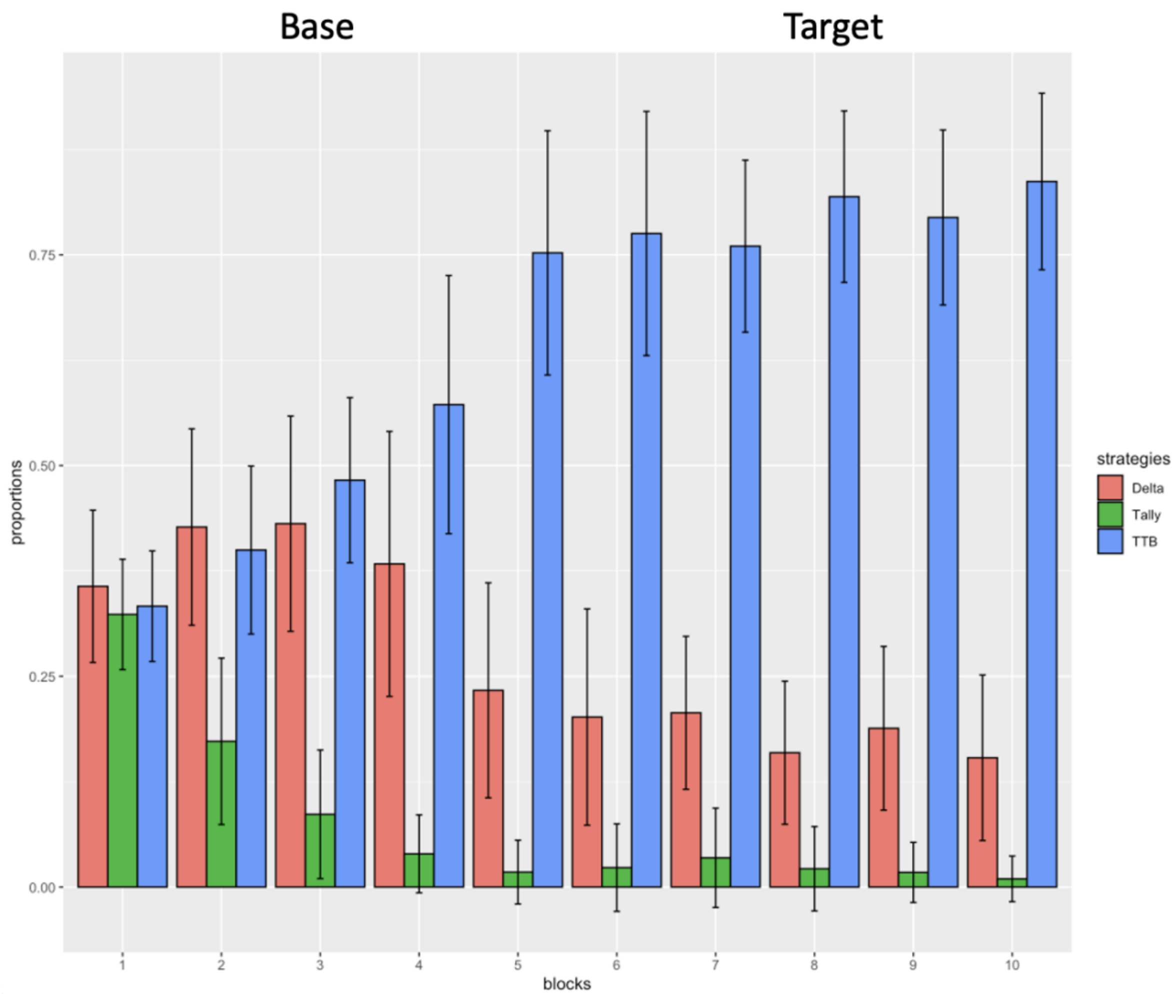

Figure 4-5 The simulation results for each block. Each bar indicates the proportion of trials in the block using the corresponding strategy. Each strategy is color coded as indicated in the legend.

We recorded the strategies applied during each trial by every simulated ACT-R agent. The proportions of each strategy's use bear similarity to the machine learning classification of behavioral data shown in Figure 3-7, broken down by block and measured by participants. The employment of the reinforced strategy, TTB, sees a steady increase for each block, while the usage of Tally and Δ-inference strategies undergo a steep decline.

Contrary to the previously noted discrepancy, the simulated classification fails to capture the drop in TTB usage during the shift to the target task, a pattern clearly evident in the human subjects' data presented in Figure 3-7. Instead, simulated agents persistently favor the best-performing strategy from the base task more frequently.

Conclusion

In this chapter, we used the ACT-R cognitive architecture to simulate the strategy selection process across two distinct environments and to model the analogical reasoning process that transpires during the transition between these environments. This cognitive framework facilitated the implementation of various heuristic strategies, and each was delineated as a sequence of steps associated with a distinct utility value. The utility value of a strategy dynamically shifts based on the feedback received by the simulated agent, thus governing the probability of selecting a particular strategy. This mechanism essentially encapsulates the learning process within the strategy selection framework.

Based on the path-mapping theory, we incorporated an existing analogical mapping method into our strategy selection process. This integration was designed to simulate the behavior of decision-makers who employ analogical reasoning in their decision-making process. In this way, we modeled the transfer of learning from one context (the base task) to another (the target task) via analogical reasoning. Using the ACT-R cognitive architecture, our study provides a rough draft of the cognitive mechanisms underlying how analogical transfer influences the decision-making processes. Moreover, it underscores the potential of cognitive architectures like ACT-R in creating sophisticated and accurate human cognition and behavior models, thereby opening avenues for further research and development in cognitive modeling, artificial intelligence, and related fields.

Moving forward, we aim to juxtapose the results derived from our integrated ACT-R model with those obtained from contemporary analogy models, such as SME. This comparison will not only foster a better understanding of the varying cognitive processes but will also aid in identifying any gaps or overlapping areas. Ultimately, our future direction entails incorporating salient features from these models into our ACT-R model, thereby enriching its predictive and explanatory power.

Chapter 5

General Discussion

Knowledge transfer, a prominent occurrence in psychological studies, considerably impacts problem-solving procedures. This knowledge encompasses cognitive skills such as strategies employed in a task, as well as structural relations concerning the task. Yet, it is often under-explored in the realm of multi-attribute decision-making, where attributes are commonly enveloped within a narrative that can, at times, trigger knowledge transfer. The research endeavors within this dissertation engage participants with two tasks that are interlinked through analogy. In this concluding chapter, I recapitulate the insights derived from this dissertation and underscore potential trajectories for future research.

What We Have Learned

The work reported in this dissertation takes the challenge proposed by Forbus et al., (2017) that encourages researchers to incorporate a model of analogical matching to various cognitive tasks. We responded to this challenge by connecting analogy with multi-attribute decision making, which are two research areas that hardly communicate.

We have shown the influence of analogical transfer on multi-attribute decision-making through the execution of a behavioral experiment, as outlined in Chapter 2. Navigating two separate tasks, the participants leveraged the first to enhance their understanding of cues and how the task worked. A dichotomy was observed, where half of our participants received an analogous scenario to their ensuing task while the remainder received a conceptually unrelated task. In the second task, however, all participants tackled the task that is either water flow or heat flow systems, equipped only with attribute names devoid of importance. The group trained with the analogous task seemed primed to apply their learned strategies to the second task. We identified the influence of analogical transfer in the analogy conditions. Participants who had previously learned the analog in the base task effectively mapped the corresponding cues to the target task. This mastery of cue relationships and cue validities conferred an advantage as participants commenced the target task equipped with more information and a higher accuracy rate. We saw improved performance, but the question was what was underlying this performance.

Moving forward to Chapter 3, we explored strategy selection in these two environments, employing a Machine Learning Strategy Identification approach (Fang et al., 2022). With the labeled behavioral data, we used machine learning models to track participants' strategies as they adapted to the environments. We first validated this machine learning method by benchmarking it against other strategy recovery techniques before applying it to our analogy experiment. By successfully pinpointing the strategy selection processes in both environments, we were able to gauge the impact of the initial environment on strategy selection in the subsequent one.

In Chapter 4, we adopted the ACT-R cognitive architecture to simulate the strategy selection process in the three environments and to model the analogical reasoning process during the transition between them. ACT-R, providing a cognitive blueprint of the processes involved in perception, memory, and decision-making, allowed us to implement various heuristic strategies as sequences of steps, each with a distinct utility value. This value, which fluctuates based on the simulated agent's feedback, ultimately governs the likelihood of a particular strategy being selected. Moreover, the existing analogical mapping method was integrated with strategy selection to simulate decision-makers behavior employing an analogy. This chapter demonstrates cognitive mechanisms underpinning analogical transfer in decision-making through the lens of ACT-R.

What We Still Need to Learn

This dissertation has shown the effects of analogy on the decision-making process and has demonstrated an ACT-R model that combines analogical mapping with decision strategies. The work has implications for information search and cue ranking mechanisms. However, these implications require further exploration to comprehend their extent and impact fully. Subsequently, I will discuss the key issues emergent from the research conducted in this dissertation and provide recommendations for potential avenues of future exploration. The main issues are the knowledge representations used in the modeling processes and the inferences they can make.

Knowledge representations

As we have illustrated the analogical mapping process using the ACT-R path mapping theory in Chapter 4, it is crucial to recognize that the model's performance depends heavily on the knowledge representations. The specific form of representation employed can significantly impact the resulting analogical mapping, making the model's performance sensitive to the representation structure. For example, when encoding a descriptive story, opting for a flat or hierarchical knowledge representation can considerably affect the analogical mapping. In the analogy model that we implemented in Chapter 4, the hierarchical structures are identical between the two domains, which results in perfect mapping. A slight change in the structure may cause failure to produce a map. This issue is also true for SME as well as other analogy models.

This sensitivity introduces the challenge of tailorability (Forbus et al., 2017), which refers to the extent that representation choices the modeler makes influence a model's results to achieve desired simulation outcomes rather than being determined by theoretical constraints. The issue of tailorability often arises when establishing correspondences between different knowledge representations, as it typically necessitates human intervention. Forbus et al. (2017) emphasize addressing this issue to ensure the validity and generalizability of modeling results. To minimize tailorability, Forbus et al. (2017) suggest using representations developed by other researchers to avoid arbitrary representation choices, which could lead to biased or unrepresentative simulation outcomes.

Not only would the knowledge representation that feeds into the model cause a great difference in the mapping result, but there is evidence showing people use different representations in constructing an analogy. Individual differences in generating analogies do exist, suggesting that unique knowledge representations can lead to different analogies. Lee and Holyoak (2008) observed that the robustness of causal models in knowledge representations affects the probability of creating accurate analogies. More robust causal models resulted in precise analogies, while weaker ones led to unrelated ones. Vendetti et al. (2015) acknowledged that students possess varying capabilities in generating and comprehending analogies, which can be attributed to factors such as prior knowledge, cognitive flexibility, and experience. The use of relational words to describe problems (e.g., "top," "middle") can further aid in the generation of analogies. Consequently, cognitive models of analogy should account for individual differences when constructing the knowledge representations to incorporate.

Another critical aspect of cognitive models is that the mapping construction process is not goal-directed. Analogists may selectively favor correspondences in the source domain that address questions in the target domain, generating corresponding candidate inferences (Gentner, 1983; Holyoak et al., 1995). This implies that mappings guided by the analogist's questions directly generate goal-relevant candidate inferences instead of relying on the unconstrained generation and assessing all possible inferences following the mapping process (Holyoak, 2012). However, many computational models of analogy, including SME (Falkenhainer et al., 1989) and ACT-R path mapping, cannot guarantee that candidate inferences will be relevant to the analogist's purpose for using the analogy, as their primary function is to identify the best structural alignment between two given representations (Forbus et al., 2017).

Given the previously mentioned issue of individual differences in knowledge representations, it is essential to develop knowledge representations that effectively capture analogists' domain knowledge for input into the model, which can then generate corresponding analogical mappings. We propose constructing these knowledge representations using semantic networks, as they can provide unique insights into the complexity of cognitive systems and the processes within them. It is based on mathematical graph theory and can quantitatively represents cognitive systems using network structures (e.g., Boccaletti et al., 2006; Chan & Vitevitch, 2009; Kenett et al., 2014; Siew et al., 2019). It represents knowledge as networks with nodes and links connecting these nodes. For example, a semantic network comprises nodes representing individual words connected if they share a semantic relationship based on co-occurrences or free associations (De Deyne et al., 2016; Steyvers & Tenenbaum, 2005). Semantic networks have been used in various research domains, including semantic processing (Kenett et al., 2017), word learning in children and adults (Goldstein & Vitevitch, 2014), and higher-order cognitive processes like creativity (Kenett et al., 2014). Representing knowledge as networks enables the application of network techniques to examine further these cognitive networks' underlying structural properties (Baronchelli et al., 2013; Borge-Holthoefer & Arenas, 2010; Morris et al., 2014). We propose that employing external representations, such as semantic networks, can facilitate understanding the analogical mapping process between individuals who utilize analogies and those who do not. This approach would help ensure that the generated inferences align more closely with the desired outcomes, minimizing the potential for tailorability. As a result, this could lead to more robust and dependable outcomes in models of analogical reasoning. We conducted a pilot study to examine the differences in the semantic networks between participants who utilized analogies and those who did not when assigned to analogically related tasks. Significant disparities were identified in the nodes they generated and the connections established between these nodes. For more details, please refer to Appendix C.

Negative transfer

The research and modeling undertaken in our study primarily demonstrated the positive influence of analogical transfer; however, it's crucial to note that analogical transfer can also yield negative impacts under certain circumstances. Typically, this negative transfer transpires when the context or rules of an initial task are in conflict with a new one, precipitating flawed assumptions or the misuse of strategies.

Consider, for instance, the international business realm where a strategy that proved successful in one country is erroneously deployed in another, oblivious to the variations in culture, economy, or legal frameworks. An effective marketing approach in a Western nation may not only fail to resonate but might even cause offense among audiences in an Eastern country, culminating in unfavorable business outcomes.

Despite their impact, convincing instances of negative transfer are comparatively scarce in decision-making literature. Singley and Anderson (1989) suggest that instances of negative transfer are commonplace in personal anecdotes, such as complaints about using a new version of a computer software program or driving in Great Britain after learning to drive in the United States. They documented negative transfer in a cognitive skill context - specifically, a text editing task. Participants who learned to use one text editor were then introduced to a new editor with different key bindings, leading to the import of nonoptimal methods from the training editor.

In a different study, Chen and Daehler (1989) examined the negative transfer of problem-solving among 6-year-olds. They found that while the children could abstract principles from stories and apply them to real problems, training that facilitated the extraction of abstract representation improved performance under positive conditions but didn't enhance their ability to discern the effectiveness of a solution under negative transfer conditions. This observation suggests that the influence of negative transfer is not only enduring but also difficult to counteract.

The frequency of concrete demonstrations of negative transfer in the decision-making literature is relatively low. Yet, the experimental design employed in this dissertation offers potential avenues to study negative transfer. This could be achieved by altering the cue validities of corresponding cues in the target task, like rendering a crucial cue from the base task unimportant in its analogous cue within the target task. The ensuing research questions then arise: What is the extent of negative transfer's impact, and how long does this effect endure? Exploring these queries could provide a more nuanced understanding of the complexities of decision-making processes.

Conclusion

The work presented in this dissertation has deepened our understanding of the impact of analogy on multi-attribute decision making, specifically in how past knowledge aids cue ranking and strategy selection in the current task. Our work explores how analogical reasoning can significantly influence strategy selection and, ultimately, the decision-making process in multi-attribute tasks. The insights gleaned not only shed light on this complex dynamic but also opened avenues for further inquiry and applications in the realm of decision science.

Appendix A

We applied MLSI to data from Walsh and Gluck (2016; henceforth referred to as W&K). The data of Walsh and Gluck (2016) is available at https://github.com/mdlee/switchingStrategies. The task in W&K is similar to our Experiment II, in which participants were given four cues and the corresponding cue validities. Similarly, in W&K, the cue values were covered by boxes, and participants needed to click on a box to see the values. Therefore, we were able to extract some of the same features from their experiment that we used in our analysis of Experiment II. We used the Random Forest model trained in our Experiment II to identify the strategies used by participants in W&K. We extracted from the W&K data twenty features, including the time needed to process the values hidden behind boxes, search orders, and strategy outcome predictions. However, there are features that our model uses that are not available in their data. For example, the decision time for each trial is not given, so we approximated this feature by summing the time needed to process each box, although the actual decision time could be greater than this sum. In addition, participants were trained to think aloud when making decisions in W&K, so the total decision times were likely longer than those in our Experiment II. Moreover, the positions of the cues were randomized for each participant in W&K, whereas cues were presented from the most to the least important for all our participants. Nevertheless, we can calculate the search order features by combining the stimulus and search information for each trial provided by W&K. Despite these differences, the ML model trained in our Experiment II identifies the strategies used by W&K's participants on a trial-by-trial basis. Figure 1 shows these identification results. Lee & Gluck (2019) developed Bayesian methods to identify when W&K's participants switched from one strategy to another. Their analysis used decision, search, and verbal report data (see Figure 14, Lee & Gluck, 2019).

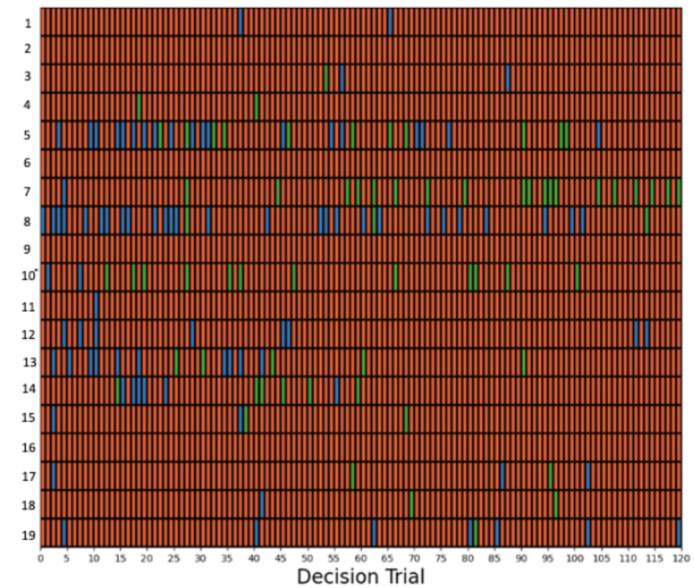

Figure 1. Strategy identification about all 120 trials for 19 participants in think aloud condition in Walsh & Gluck (2016) using trained ML model from the current study. Red = TTB; green = Tallying; Blue = WADD.

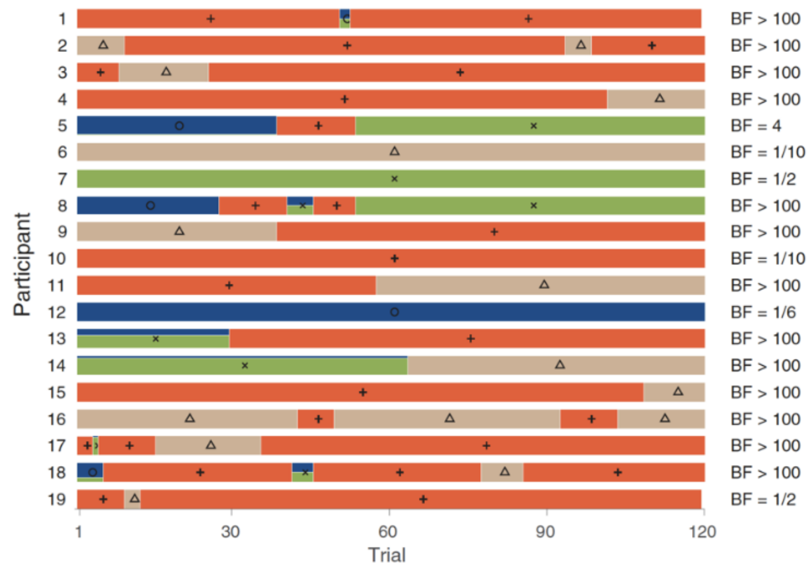

Figure 14 from Lee & Gluck's (2019). Inferences about switching strategy use for all 19 participants based on decision, search, and verbal report data. Red = TTB (+ marker); green = tally (X marker); blue = WADD (circle marker); beige = guess (triangle marker).

Note that because guessing is not a candidate strategy in our study, the trained Random Forests model cannot identify a participant using guessing. Both methods identify some interesting cases that would be missed by methods that assume participants use the same strategy throughout. For example, both our Random Forest model and the Lee & Gluck's method indicate that Participant 1 used TTB most of the time but briefly adopted to a compensatory strategy for some trials midway through the experiment, and Participants 5, 8, 13 and 18 used a mixture of WADD, TTB, and Tallying. There are also agreements on predicting TTB. For Participants 2, 3, 4, 9, 10, 11, 15, 17, and 19 most of the trials are identified as using TTB by both approaches. The ML model predicts the trials that were identified as guessing in the Bayesian method as using TTB, because decision time and search behaviors of guessing are more similar to TTB than the other two strategies in the data. No doubt, the comparison between the two approaches would be more informative if the approaches had been applied to data from an experiment designed for the purpose of comparing the approaches. Such a systematic comparison, and perhaps integration, of the two approaches could be an interesting project for future investigation.

Appendix B-1

The flowchart of implementing the Δ-inference model in ACT-R

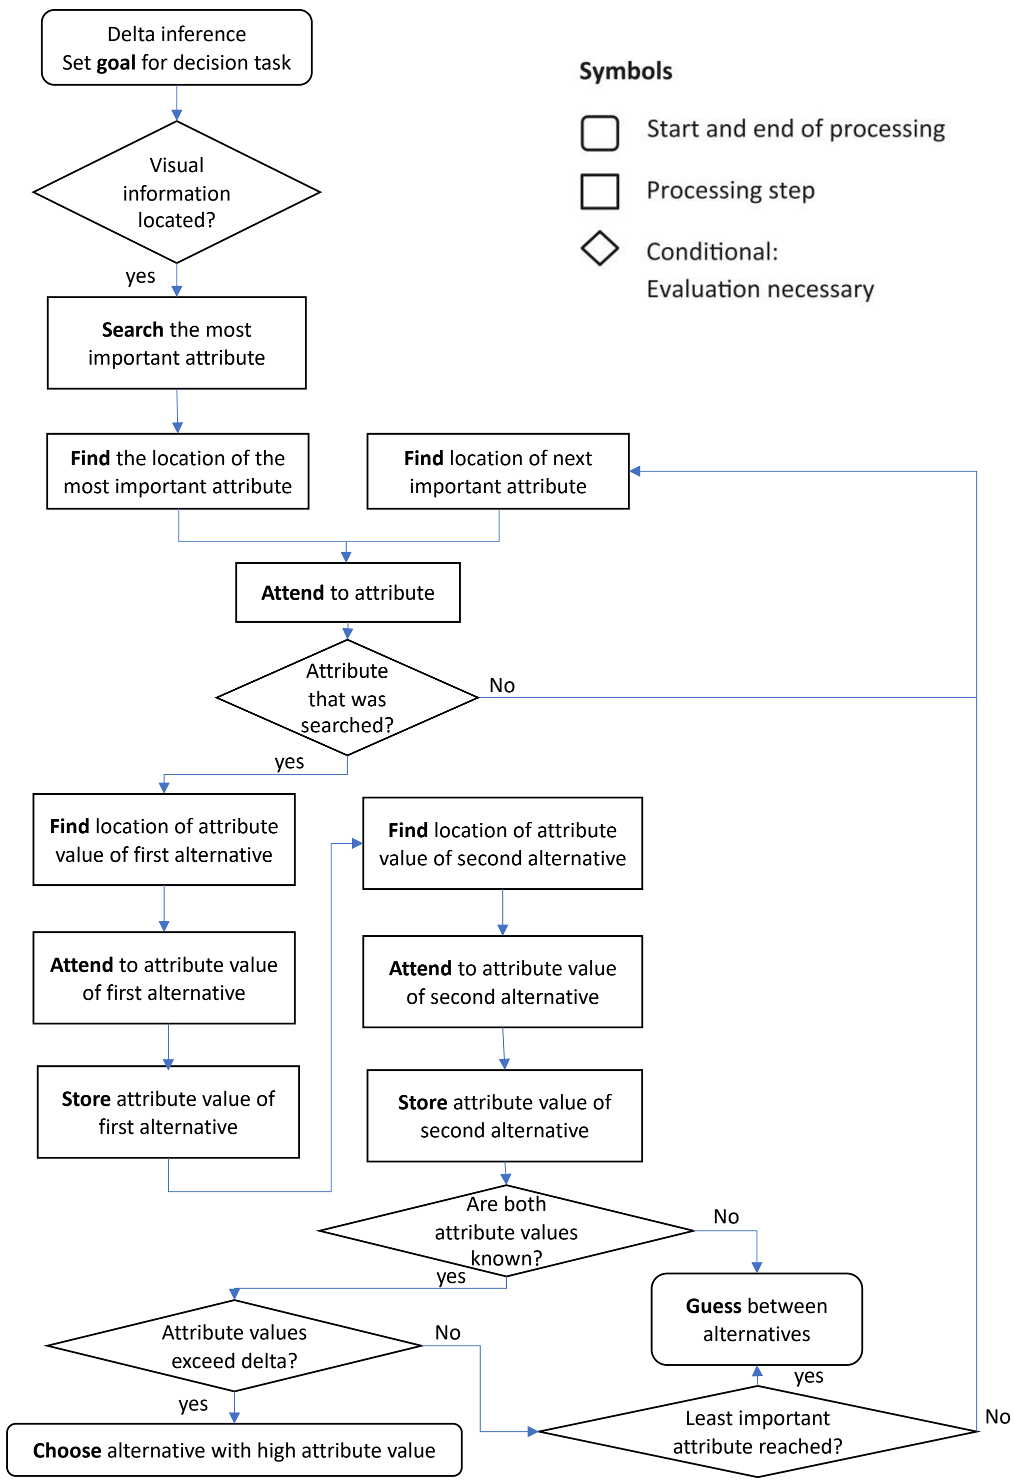

Appendix B-2

The flowchart of implementing the Tallying model in ACT-R

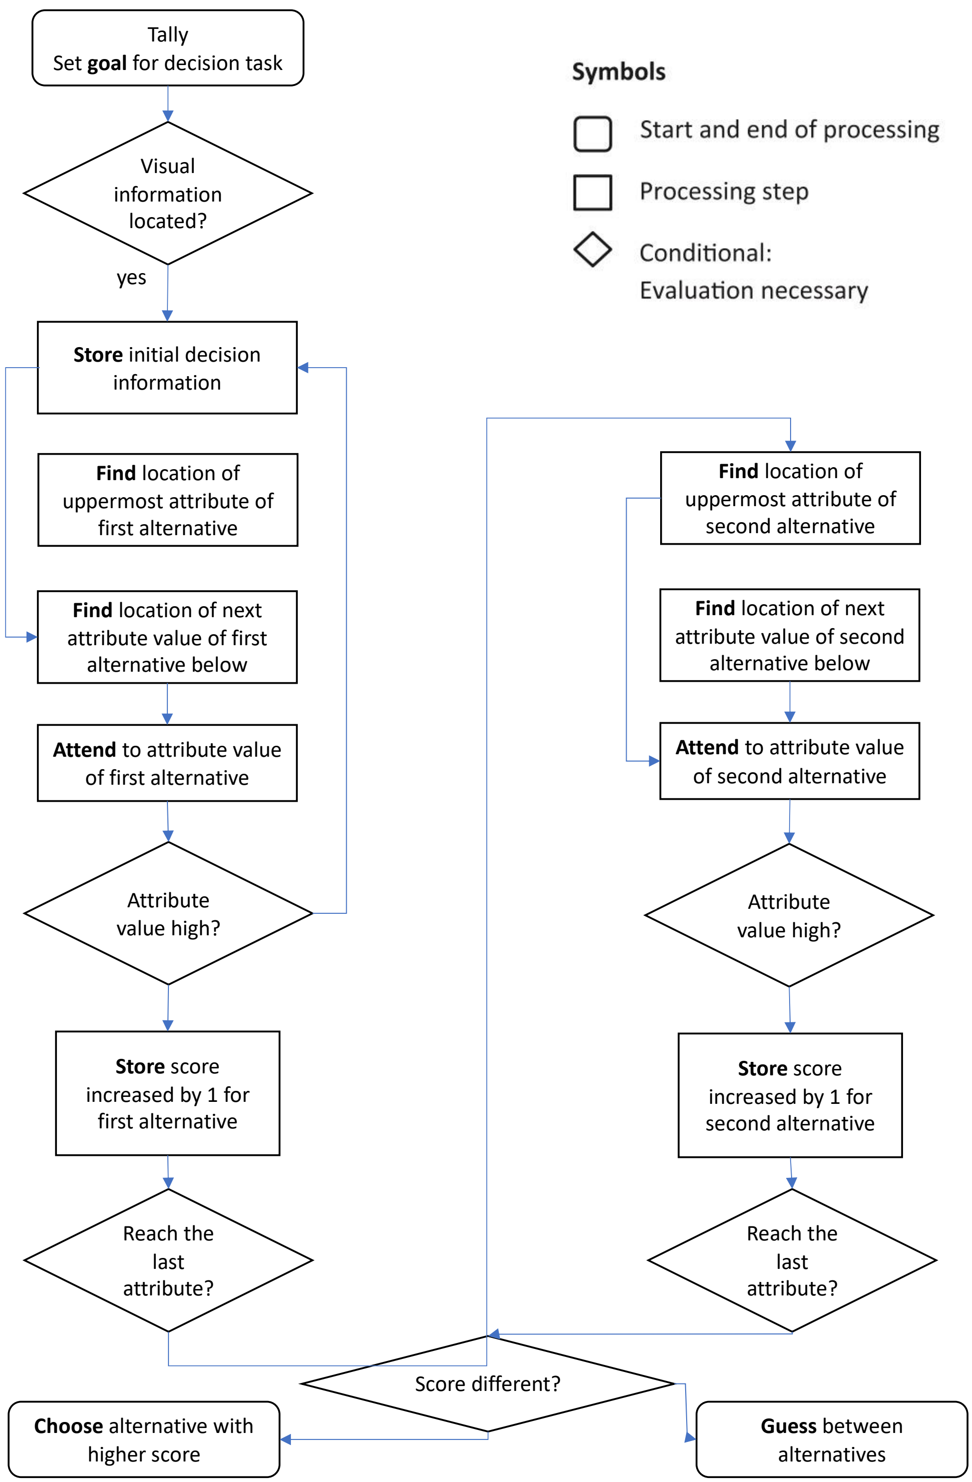

Appendix C

Participants were given the same water and heat flow tasks described in Chapter 2. In addition to the experiment, they were required to do a free recall task for 2 minutes at the beginning of the base and target task. A survey containing cue mapping was given at the end. Twenty participants were recruited for this study.

Experiment Design

Measure the knowledge representation

Words Free recall 1

Words Free recall 2

Conditions	Pre-test (no feedback)	Base task (with cue validity)	Pre-test (no feedback)	Target task (no cue validity)	Survey
Analogy condition		Water flow		Heat flow	cue mappings

Participants were given a word (e.g., water flow, heat flow) and free recall words for 2 minutes

We identified participants using analogy vs. not using analogy by the end survey. We found that 12 participants used analogy, who gave perfect cue mapping in the survey, and 8 participants did not.

Analogy group answer example:
Which cue in Task 1 is most similar to the bar diameter in Task 2? Pipe diameter
…
Does information learned in Task 1 help you to make decisions in Task 2?
Yes, it gave me practice and it told me how important each cue was when it came to water flow, so I applied the same idea to heat flow.

Base: Water flow semantic network

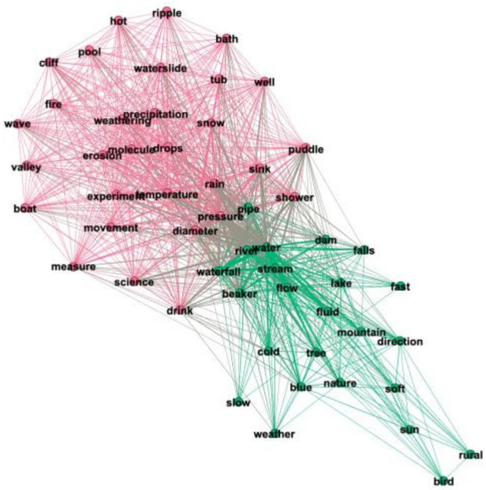

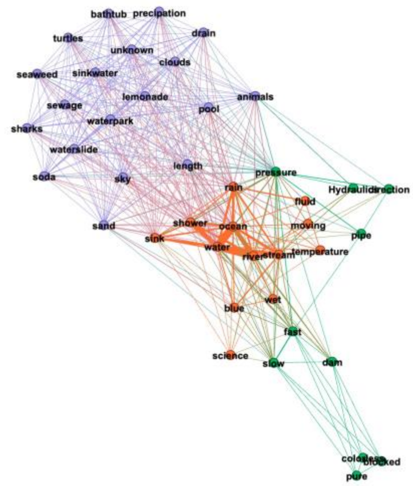

Group	Analogy	No analogy
# of nodes	56	42
Clustering Coefficient	0.868	0.849

Target: Heat flow semantic network

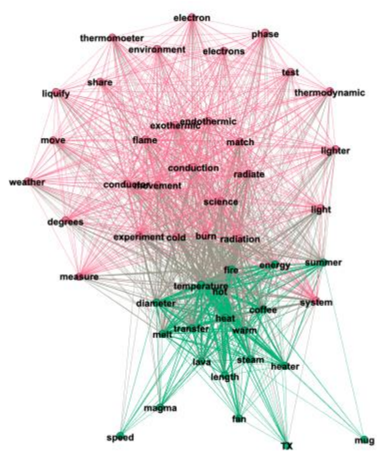

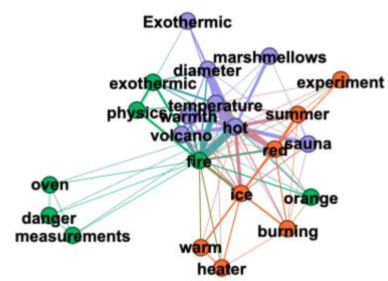

Group	Analogy	No analogy
# nodes	48	22
Clustering Coefficient	0.91	0.807

References

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological review*, *98*(3), 409.

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological science*, *2*(6), 396-408.

Anderson, J. R., Bothell, D., Lebiere, C., & Matessa, M. (1998). An integrated theory of list memory. *Journal of Memory and Language, 38*(4), 341-380.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological review*, *111*(4), 1036.

Anderson, J. R. (2009). How can the human mind occur in the physical universe?. Oxford University Press.

Anderson, J. R., Betts, S. A., Ferris, J. L., & Fincham, J. M (2011). Using Brain imaging to interpret student problem solving. In *IEEE Intelligent Systems*, 26(5), 22-29.

Anderson, J. R., Betts, S., Bothell, D., & Lebiere, C. (2021). Discovering Skill. Cognitive Psychology, 129.

Archer, K. J., & Kimes, R. V. (2008). Empirical characterization of random forest variable importance measures. Computational Statistics & Data Analysis, 52, 2249–2260.

Ashby, F. G., & Maddox, W. T. (1992). Complex decision rules in categorization: contrasting novice and experienced performance. *Journal of Experimental Psychology: Human Perception and Performance*, *18*(1), 50.

Baronchelli, A., Ferrer-i-Cancho, R., Pastor-Satorras, R., Chater, N., & Christiansen, M. H. (2013). Networks in cognitive science. *Trends in cognitive sciences*, *17*(7), 348-360.

Beach, L. R., & Mitchell, T. R. (1978). A contingency model for the selection of decision strategies. *Academy of management review*, *3*(3), 439-449.

Bettman, J. R., Johnson, E. J., & Payne, J. W. (1990). A componential analysis of cognitive effort in choice. Organizational Behavior and Human Decision Processes, 45, 111–139.

Boccaletti, S., Latora, V., Moreno, Y., Chavez, M., & Hwang, D. U. (2006). Complex networks: Structure and dynamics. Physics reports, 424(4-5), 175-308.

Borst, J. P., & Anderson, J. R. (2013). Using model-based functional MRI to locate working memory updates and declarative memory retrievals in the fronto-parietal network. *Proceedings of the National Academy of Sciences U.S.A., 110*(5), 1628-1633.

Borge-Holthoefer, J., & Arenas, A. (2010). Semantic networks: Structure and dynamics. *Entropy*, *12*(5), 1264-1302.

Bobadilla-Suarez, S., & Love, B. C. (2018). Fast or frugal, but not both: Decision heuristics under time pressure. Journal of Experimental Psychology: Learning, Memory, and Cognition, 44(1), 24.

Bröder, A. (2000). Assessing the empirical validity of the ‘Take-the-best’ heuristic as a model of human probabilistic inference. Journal of Experimental Psychology: Learning, Memory, and Cognition, 26, 1332–1346.

Bröder, A. (2003). Decision making with the "adaptive toolbox": Influence of environmental structure, intelligence, and working memory load. Journal of experimental psychology. Learning, memory, and cognition, 29 (4), 611.

Busemeyer, J. R., & Myung, I. J. (1992). An adaptive approach to human decision making: Learning theory, decision theory, and human performance. *Journal of Experimental Psychology: General*, *121*(2), 177.

Bush, R. R., & Mosteller, F. (1955). Stochastic models for learning.

Brighton, H., & Gigerenzer, G. (2012). Are rational actor models “rational” outside small worlds. *Evolution and rationality: Decisions, co-operation, and strategic behavior*, 84-109.

Brown, A. L., & Kane, M. J. (1988). Preschool children can learn to transfer: Learning to learn and learning from example. *Cognitive psychology*, *20*(4), 493-523.

Bröder, A., & Gaissmaier, W. (2007). Sequential processing of cues in memory-based multiattribute decisions. *Psychonomic Bulletin & Review*, *14*(5), 895-900.

Brusovansky, M., Glickman, M., & Usher, M. (2018). Fast and effective: Intuitive processes in complex decisions. *Psychonomic Bulletin & Review*, *25*, 1542-1548.

Canellas, M. C., & Feigh, K. M. (2017). Heuristic information acquisition and restriction rules for decision support. *IEEE Transactions on Human-Machine Systems*, *47*(6), 939-950.

Canini, K. R., Shashkov, M. M., & Griffiths, T. L. (2010, June). Modeling Transfer Learning in Human Categorization with the Hierarchical Dirichlet Process. In *ICML* (pp. 151-158).

Chan, K. Y., & Vitevitch, M. S. (2009). The influence of the phonological neighborhood clustering coefficient on spoken word recognition. Journal of Experimental Psychology: Human Perception and Performance, 35(6), 1934.

Chen, Z., & Daehler, M. W. (1989). Positive and negative transfer in analogical problem solving by 6-year-old children. *Cognitive Development*, *4*(4), 327-344.

Christensen-Szalanski, J. J. (1978). Problem solving strategies: A selection mechanism, some implications, and some data. *Organizational Behavior and Human Performance*, *22*(2), 307-323.

Costa-Gomes, M., Crawford, V. P., & Broseta, B. (2001). Cognition and behavior in normal-form games: An experimental study. Econometrica, 69(5), 1193-1235.

Davis-Stober, C. P., & Brown, N. (2011). A shift in strategy or" error"? Strategy classification over multiple stochastic specifications. Judgment and Decision Making, 6(8), 800-813.

De Deyne, S., Kenett, Y. N., Anaki, D., Faust, M., & Navarro, D. (2017). Large-scale network representations of semantics in the mental lexicon.

Dougherty, M. R., Franco-Watkins, A. M., & Thomas, R. (2008). Psychological plausibility of the theory of probabilistic mental models and the fast and frugal heuristics. *Psychological review*, *115*(1), 199.

Dimov, C.M., Marewski, J. N., & Schooler, L. J. (2017). Architectural process models of decision making: Towards a model database. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. J. Davelaar (Eds.), Proceedings of the 39th Annual Conference of the Cognitive Science Society (pp. 1931-1936). Austin, TX: Cognitive Science Society.

Erev, I., & Roth, A. E. (2001). Simple reinforcement learning models and reciprocation in the prisoner's dilemma game.

Fang, J., Schooler, L., & Shenghua, L. (2022). Machine learning strategy identification: A paradigm to uncover decision strategies with high fidelity. *Behavior Research Methods*, *55*(1), 263-284.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial intelligence*, *41*(1), 1-63.

Fechner, H. B., Pachur, T., Schooler, L. J., Mehlhorn, K., Battal, C., Volz, K. G., & Borst, J. P. (2016). Strategies for memory-based decision making: Modeling behavioral and neural signatures within a cognitive architecture. *Cognition*, *157*, 77-99.

Fechner, H. B., Schooler, L. J., & Pachur, T. (2018). Cognitive costs of decision-making strategies: A resource demand decomposition analysis with a cognitive architecture. *Cognition*, *170*, 102-122.

Forbus, K. D., Ferguson, R. W., Lovett, A., & Gentner, D. (2017). Extending SME to handle large-scale cognitive modeling. *Cognitive Science*, *41*(5), 1152-1201.

Friedman, J., Hastie, T., & Tibshirani, R. (2001). The elements of statistical learning (Vol. 1, No. 10). New York: Springer series in statistics.

Fu, W. T., & Anderson, J. R. (2006). From recurrent choice to skill learning: a reinforcement-learning model. *Journal of experimental psychology: General*, *135*(2), 184.

Gavetti, G., Levinthal, D. A., & Rivkin, J. W. (2005). Strategy making in novel and complex worlds: The power of analogy. *Strategic Management Journal*, *26*(8), 691-712.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive science*, *7*(2), 155-170.

Gentner, D., Loewenstein, J., & Thompson, L. (2003). Learning and Transfer: A General Role for Analogical Encoding. *Journal of Educational Psychology*, *95*, 393-408.

Gentner, D. (2010). Bootstrapping the mind: Analogical processes and symbol systems. *Cognitive Science*, *34*(5), 752-775.

Gentner, D. (2017). Analogy. *A companion to cognitive science*, 107-113.

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual review of psychology*, *62*, 451-482.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. Psychological Review, 103(4), 650-669.

Gigerenzer, G., & Todd, P. M. (1999). ABC Research Group. *Simple heuristics that make us smart*, 3-34.

Gigerenzer, G., & Selten, R. (Eds.). (2002). Bounded rationality: The adaptive toolbox. MIT press.

Goldstein, R., & Vitevitch, M. S. (2014). The influence of clustering coefficient on word-learning: How groups of similar sounding words facilitate acquisition. *Frontiers in psychology*, *5*, 1307.

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, *27*(4), 591-635.

Glöckner, A. (2009). Investigating intuitive and deliberate processes statistically: The multiple-measure maximum likelihood strategy classification method. Judgment and Decision Making, 4(3), 186-199.

Harte, J. M., & Koele, P. (2001). Modelling and describing human judgement processes: The multiattribute evaluation case. Thinking & reasoning, 7(1), 29-49.

Hertwig, R. E., & Hoffrage, U. E. (2013). *Simple heuristics in a social world*. Oxford University Press.

Hilbig, B. E., & Moshagen, M. (2014). Generalized outcome-based strategy classification: Comparing deterministic and probabilistic choice models. Psychonomic Bulletin & Review, 21, 1431–1443.

Holyoak, K. J., Thagard, P., & Sutherland, S. (1995). Mental leaps: analogy in creative thought. Nature, 373(6515), 572-572.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive science*, *13*(3), 295-355.

Holyoak, K. J. (2012). Analogy and relational reasoning.

Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological review*, *110*(2), 220.

Jekel, M., Nicklisch, A., & Glöckner, A. (2010). Implementation of the Multiple-Measure Maximum Likelihood strategy classification method in R: Addendum to Glöckner (2009) and practical guide for application. Judgment and Decision Making, 5(1), 54-63.

Johnson, E. J., Schulte-Mecklenbeck, M., & Willemsen, M. C. (2008). Process models deserve process data: Comment on Brandstätter, Gigerenzer, and Hertwig (2006). Psychological Review, 115, 263-272.

Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: individual differences in working memory. *Psychological review*, *99*(1), 122.

Katz, I., Bereby-Meyer, Y., Assor, A., & Danziger, S. (2010). Children's adaptive pre-decisional search behavior: Effects of memory and number of alternatives. *Journal of Economic Psychology*, *31*(1), 17-24.

Kenett, Y. N., Anaki, D., & Faust, M. (2014). Investigating the structure of semantic networks in low and high creative persons. *Frontiers in human neuroscience*, *8*, 407.

Kenett, Y. N., Levi, E., Anaki, D., & Faust, M. (2017). The semantic distance task: Quantifying semantic distance with semantic network path length. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *43*(9), 1470.

Kotsiantis, S. B., Zaharakis, I., & Pintelas, P. (2007). Supervised machine learning: A review of classification techniques. Emerging artificial intelligence applications in computer engineering, 160(1), 3-24.

Krefeld-Schwalb, A., Donkin, C., Newell, B. R., & Scheibehenne, B. (2019). Empirical comparison of the adjustable spanner and the adaptive toolbox models of choice. Journal of Experimental Psychology: Learning, Memory, and Cognition, 45(7), 1151.

Krol, M., & Krol, M. (2017). A novel approach to studying strategic decisions with eye-tracking and machine learning. Judgment and Decision Making, 12(6), 596-609.

Lee, H. S., & Holyoak, K. J. (2008). The role of causal models in analogical inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(5), 1111.

Lee, H.S., Betts, S. & Anderson,J.R. (2016) Learning Problem-Solving Rules as Search Through a Hypothesis Space. Cognitive Science, 40(5), 1036-1079.

Lee, M. D. (2011). How cognitive modeling can benefit from hierarchical Bayesian models. Journal of Mathematical Psychology, 55, 1–7.

Lee, M. D., & Wagenmakers, E. J. (2014). Bayesian cognitive modeling: A practical course. Cambridge university press.

Lee, M. D. (2016). Bayesian outcome-based strategy classification. Behavior Research Methods, 48, 29–41.

Lee, M. D., Gluck, K. A., & Walsh, M. M. (2019). Understanding the complexity of simple decisions: Modeling multiple behaviors and switching strategies. Decision, 6(4), 335.

Lee, M. D., & Gluck, K. A. (2020). Modeling Strategy Switches in Multi-attribute Decision Making. Computational Brain & Behavior, 1-16.

Lieder, F., & Griffiths, T. L. (2017). Strategy selection as rational metareasoning. Psychological Review, 124(6), 762–794.

Lovett, M. C., & Anderson, J. R. (1996). History of success and current context in problem solving: Combined influences on operator selection. *Cognitive psychology*, *31*(2), 168-217.

Marewski, J. N., Gaissmaier, W., & Gigerenzer, G. (2010). Good judgments do not require complex cognition. *Cognitive processing*, *11*, 103-121.

Marewski, J. N., & Schooler, L. J. (2011). Cognitive niches: An ecological model of strategy selection. *Psychological review*, *118*(3), 393.

Marewski, J. N., & Link, D. (2014). Strategy selection: An introduction to the modeling challenge. Wiley Interdisciplinary Reviews: Cognitive Science, 5, 39-59. DOI: 10.1002/wcs.1265

Meyer, D. E., & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part I. Basic mechanisms. *Psychological review*, *104*(1), 3.

Morais, A. S., Schooler, L., Olsson, H., & Meder, B. (2014). From Causal Models to Sound Heuristic Inference. In *CogSci*.

Morey, R. D., Rouder, J. N., Jamil, T., & Morey, M. R. D. (2015). Package 'bayesfactor'. *URLh http://cran/r-projectorg/web/packages/BayesFactor/BayesFactor pdf i (accessed 1006 15)*.
Nellen, S. (2003). The use of the "take-the-best" heuristic under different conditions, modelled with ACT-R. In Proceedings of the (pp. 171-76).

Nelson, J. D., Divjak, B., Gudmundsdottir, G., Martignon, L. F., & Meder, B. (2014). Children's sequential information search is sensitive to environmental probabilities. *Cognition*, *130*(1), 74-80.

Newell, A. (1994). *Unified theories of cognition*. Harvard University Press.

Newell, B. R., Collins, P., & Lee, M. D. (2007). Adjusting the spanner: Testing an evidence accumulation model of decision making. In Proceedings of the Annual Meeting of the Cognitive Science Society (Vol. 29, No. 29).

Oravecz, Z., & Muth, C. (2018). Fitting growth curve models in the Bayesian framework. *Psychonomic Bulletin & Review*, *25*, 235-255.

Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, *22*(10), 1345-1359.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of experimental psychology: Learning, Memory, and Cognition*, *14*(3), 534.

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

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... & Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. Journal of Machine Learning Research, 12, 2825-2830.

Raftery, A. E. (1995). Bayesian model selection in social research. *Sociological methodology*, 111-163.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations on the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. R. Prokasy (eds.), *Classical Conditioning: II. Current Research and Theory* (pp. 64-99). New York: Appleton-Century-Crofts.

Rieskamp, J., & Otto, P. E. (2006). SSL: a theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*(2), 207.

Rieskamp, J., & Hoffrage, U. (2008). Inferences under time pressure: How opportunity costs affect strategy selection. Acta Psychologica, 127, 258–276.

Richland, L. E., & Simms, N. (2015). Analogy, higher order thinking, and education. *Wiley Interdisciplinary Reviews: Cognitive Science*, *6*(2), 177-192.

Riedl, R., Brandstätter, E., & Roithmayr, F. (2008). Identifying decision strategies: A process- and outcome-based classification method. Behavior research methods, 40(3), 795-807.

Salvucci, D. D., & Anderson, J. R. (2001). Integrating analogical mapping and general problem solving: the path-mapping theory. *Cognitive Science*, *25*(1), 67-110.

Salvucci, D. D., & Taatgen, N. A. (2008). Threaded cognition: an integrated theory of concurrent multitasking. *Psychological review*, *115*(1), 101.

Scheibehenne, B., Rieskamp, J., & Wagenmakers, E.-J. (2013). Testing adaptive toolbox models: A Bayesian hierarchical approach. Psychological Review, 120, 39–64.

Schooler, L. J., & Hertwig, R. (2005). How forgetting aids heuristic inference. *Psychological review*, *112*(3), 610.

Schulte-Mecklenbeck, M., Kühberger, A., Gagl, B., & Hutzler, F. (2017). Inducing thought processes: Bringing process measures and cognitive processes closer together. Journal of Behavioral Decision Making, 30(5), 1001-1013.

Siew, C. S., Wulff, D. U., Beckage, N. M., & Kenett, Y. N. (2019). Cognitive network science: A review of research on cognition through the lens of network representations, processes, and dynamics. *Complexity*, *2019*.

Singley, M. K., & Anderson, J. R. (1989). *The transfer of cognitive skill* (No. 9). Harvard University Press.

Stewart, T. R. (1988). Judgment analysis: procedures. In *Advances in psychology* (Vol. 54, pp. 41-74). North-Holland.

Steyvers, M., & Tenenbaum, J. B. (2005). The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. Cognitive science, 29(1), 41-78.

Stewart, T. R. (1988). Judgment analysis: procedures. In Advances in psychology (Vol. 54, pp. 41-74). North-Holland.

Taatgen, N. A. (2013). The nature and transfer of cognitive skills. *Psychological review*, *120*(3), 439.

Team, R. (2015). RStudio: integrated development for R. *(No Title)*.

Todd, P., & Dieckmann, A. (2004). Heuristics for ordering cue search in decision making. *Advances in neural information processing systems*, *17*.

Todd, P. M., & Gigerenzer, G. (2000). Précis of simple heuristics that make us smart. *Behavioral and brain sciences*, *23*(5), 727-741.

Thorndike, E. L., & Woodworth, R. S. (1901). The influence of improvement in one mental function upon the efficiency of other functions: III. Functions involving attention, observation and discrimination. *Psychological Review*, *8*(6), 553.

Van Ravenzwaaij, D., Moore, C. P., Lee, M. D., & Newell, B. R. (2014). A hierarchical Bayesian modeling approach to searching and stopping in multi-attribute judgment. Cognitive Science, 38(7), 1384-1405.

Vendetti, M. S., Matlen, B. J., Richland, L. E., & Bunge, S. A. (2015). Analogical reasoning in the classroom: Insights from cognitive science. *Mind, Brain, and Education*, *9*(2), 100-106.

Department of Psychology
Syracuse University
522 Huntington Hall, Syracuse, NY, 13210
Email: jfang100@syr.edu

EDUCATION

Ph. D., Cognitive Psychology — June 2023
Syracuse University, Syracuse, NY, USA
Advisor: Dr. Lael J. Schooler
Dissertation: *Knowledge transfer in multi-attribute decision making*

M.S., Applied Statistics — Aug 2019
Syracuse University, Syracuse, NY, USA

M.S., Experimental Psychology — Aug 2017
Syracuse University, Syracuse, NY, USA
Advisor: Dr. Lael J. Schooler
Thesis: *Inferring decision strategies based on the path to a choice*

M.S., Economics — May 2015
Iowa State University, Ames, IA, USA
Advisor: Dr. Elizabeth Hoffman
Thesis: *The impact of feedback on energy conservation*

B.S., Economics & Mathematics — Aug 2013
Iowa State University, Ames, IA, USA
Magna cum laude

B.S., Finance — Aug 2013
Iowa State University, Ames, IA, USA
Magna cum laude & Honors program

PAPER AND CONFERENCE PRESENTATIONS

Fang, J., Schooler L.J., & Luan S. (2022). Machine Learning Strategy Identification: A Paradigm to Uncover Decision Strategies with High Fidelity. Behavior Research Methods. 1-22

Zhang, Y., Fang, J., & Krishnakumar, A., (under review). Reconceptualizing and understanding coercive control: Natural Language Process of women's help seeking at IPV online forum. Violence Against Women.

Tan M, Newman L.S., Zhang B., & Fang J. (under review). Mnemic Neglect in the United States and China: A Cross-cultural Investigation of Spontaneous Self-protection.

Fang, J., & Schooler, L. (2022, July). Analogical transfer in multi-attribute decision making. Paper presented at Virtual MathPsych/ICCM 2022. Via mathpsych.org/presentation/902.

Fang, J. & Schooler L. (2019). Knowledge transfer via Analogy in Multi-attribute decision making. Poster presented at 2019 Annual Meeting of the Society for Mathematical Psychology, Montreal Canada.

Fang, J. & Schooler L. (2018). Inferring decision strategies based on the path to a choice. Poster presented at 2018 Annual Meeting of the Society for Mathematical Psychology, Madison.

Fang, J. (2018). Supervised machine learning of mouse movements. Poster presented at Deep Learning & Reinforcement Learning Summer School 2018, Toronto, Canada.

Fang, J. (2017). Inferring decision strategies based on the path to a choice. Best poster award at the Summer Institute on Bounded Rationality at the Max Planck Institute for Human Development.

Fang, J. & Schooler L. (2016). Inferring decision strategies based on the path to a choice. Poster presented at the November 2016, Psychonomics Society's Annual meeting, Boston, USA.

Fang, J. & Lapan H. (2013). Population aging in Japan and its Implications. Poster presented at the September, 2013 Honors Project presentations, Ames, IA.

TEACHING EXPERIENCE

Iowa State University

Econ 101: Intro to Microeconomics (Spring 2013).
Econ 102: Intro to Macroeconomics (Fall 2013).
Econ 380: Money and Banking (Spring 2014)
Econ 480: Game Theory (Fall 2014).

Syracuse University

Psych 322: Cognitive Psychology (Lecturer, Summer 2019, Summer 2020)
Psych 313: Research Design and Methodology (Recitation instructor, Fall 2015, Spring 2016, Fall 2019, Spring & Fall 2020).
Psych 400: Thinking and Decision Making (Spring 2017). Guest Lecture on the Fast and Frugal Trees
Psych 205: Foundations of Human Behavior (Recitation instructor, 2016 – 2019)

AWARDS/MEMBERSHIPS

The Outstanding TA Award 2020, Syracuse University
Best Poster Reward at the Summer Institute on Bounded Rationality 2017 Member, WiSE (Women in Science and Engineer)
Member, Psychonomics Society, MathPsych
Data Science Nanodegree, Udacity
Data Structures & Algorithms Nanodegree, Udacity
Front-End Web Developer Nanodegree, Udacity
Iowa State University Dean's List all semesters
Society of Actuaries/CAS/CIA, Passed FM/2, P/1

RESEARCH INTERESTS

- Judgement and Decision making
- Heuristics
- Knowledge transfer, Analogical Reasoning
- Machine learning/Deep learning/NLP
- ACT-R Modeling