The Role of Spatial Structure in Problem Solving

Analysis at an information type level of abstraction

Rasmus Wienemann

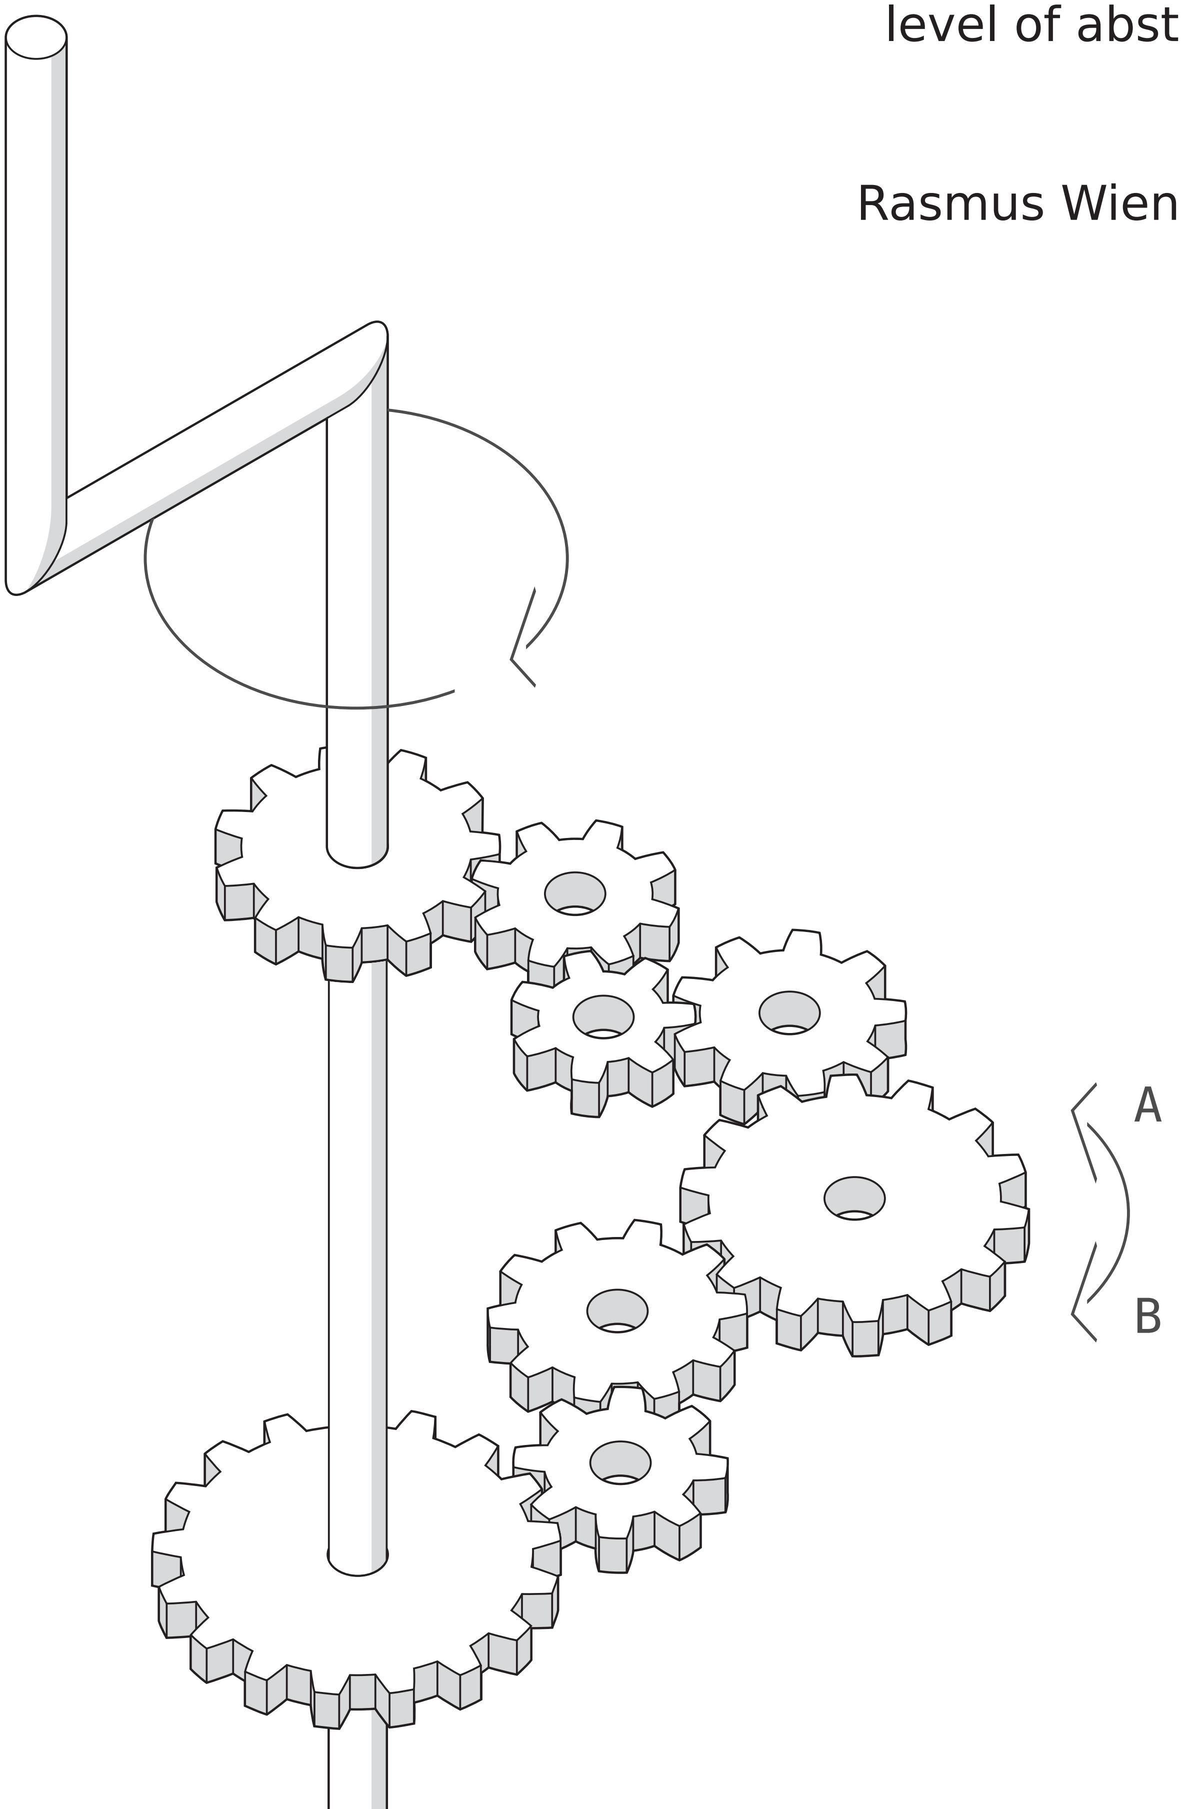

The Role of Spatial Structure in Problem Solving

Analysis at an information type level of abstraction

Rasmus Wienemann

Dissertation zur Erlangung des akademischen Grades Dr. ing.
Vorgelegt im Fachbereich 3 – Mathematik & Informatik,
Universität Bremen

Supervisors:
PD Dr.rer.nat. habil. Holger Schultheis
Prof. Dr. rer. nat. Sven Bertel

Date of colloquium:
5.10.2021

Für Marloes

Abstract

Everyday life entails a multitude of situations that can be described as (a series of) problem solving tasks. These are situated in an information rich environment from which problem relevant information needs to be gathered. Much of the interaction with the environment is bounded by its spatial structure. Not surprisingly humans have adapted to be highly capable in extracting and using spatial information from their environment. Past research suggest the existence of special-purpose processing mechanisms in human cognition for spatial information. Less is known about how humans utilize these efficient mechanisms to also aid problem solving with non spatial information.

Mental representations differ in resolution depending on the scope of what is represented. A bend in the road might be represented when imagining the route to the bakery, but not for a trip on the highway. I argue that the kind of information retrieved by problem solvers also differs in granularity depending on the applied level of abstraction. I hypothesize that this would enable one to study problem solving behavior with any information type (e.g. color, pitch, distance, topology) granularity.

Artificial intelligence systems outperform humans in extremely information rich and noisy visuo-spatial domains as real time computer games or autonomous driving. Most of these systems, though, are based on neural networks and learning algorithms that produce black boxes as to what is being analyzed and represented. The information type granularity approach lends itself to build computational models that provide such insight. We have to build cognitive agents whose behavior we understand if we want them to represent an use information types the same as people do.

In order to study problem solving at an information type granularity, tasks sensitive to different information types are needed. In this thesis I define the requirements for such experimental paradigms for studying spatial information use in problem solving. I created three experimental designs to apply the information type granularity approach to several questions about human problem solving.

1. An analogy making task to test whether humans are able to apply their spatial capabilities to novel, non-spatial domains.

2. A tic-tac-toe isomorph to determine if problem solvers seek spatial information to aid in a non-spatial domain. And if so would spatial information be detrimental to the task impede performance or be ignored?

3. A card-sorting task to verify that problem solvers don't simply prefer spatial information types over non-spatial types when competing both are equally valid to solve the task.

I further build a computational model to exemplify how my experimental paradigms could be used to test and develop cognitive agents.

In this work I show that humans are able to use spatial information in non-spatial tasks if it aids them in solving a task. I further show that this is not merely an effect of spatial information being more salient or preferred. In a setting where spatial and non-spatial information types were equal with regards to solving a problem neither was consistently preferred. Problem solvers readily ignored spatial information when it was conflicting with respect to the problem's structure. I provide a proof of concept that computational cognitive modeling applied to information type granularity processing have merit in studying human problem solving. This approach facilitates building systems that represent information in such a way that it aligns with human representations. It can be applied to create feedback systems that present information in a modality that the user can easily integrate in their mental model improving understanding.

Zusammenfassung

Das alltägliche Leben bringt eine Vielzahl von Situationen mit sich, die als (eine Reihe von) Problemlösungsaufgaben beschrieben werden können. Diese sind in einer informationsreichen Umgebung situiert, aus der problemrelevante Informationen gewonnen werden müssen. Ein Großteil der Interaktion mit der Umwelt ist durch deren räumliche Struktur begrenzt. Es überrascht daher nicht, dass Menschen eine hohe Fähigkeit entwickelt haben, um Rauminformationen aus ihrer Umwelt zu extrahieren und zu nutzen. Frühere Studien legen nahe, dass die menschliche Kognition über spezielle Verarbeitungsmechanismen für räumliche Informationen verfügt. Weniger ist darüber bekannt, wie Menschen diese effizienten Mechanismen nutzen, um das Problemlösen mit nicht räumlichen Informationen zu unterstützen.

Mentale Repräsentationen unterscheiden sich in ihrer Auflösung je nach dem Umfang dessen, was repräsentiert wird. Eine Kurve in der Straße mag repräsentiert werden, wenn man sich den Weg zur Bäckerei vorstellt, aber nicht für eine Fahrt auf der Autobahn. Ich argumentiere, dass die Art der Informationen, die von Problemlösern repräsentiert werden, sich auch in der Granularität unterscheiden. Ich stelle die Hypothese auf, dass man damit das Problemlösungsverhalten auch auf der Granularität der Informationsart (z.B. Farbe, Tonhöhe, Entfernung, Topologie) untersuchen kann.

Systeme der künstlichen Intelligenz übertreffen den Menschen in extrem informationsreichen und verrauschten visuell-räumlichen Domänen wie Echtzeit-Computerspielen oder dem autonomen Fahren. Die meisten dieser Systeme basieren jedoch auf neuronalen Netzen und Lernalgorithmen, welche Black Boxes hinsichtlich dessen, was analysiert und dargestellt wird, produzieren. Der Ansatz der Informationsgranularität bietet sich an, um Computermodelle zu kreieren, die solche Einblicke liefern können. Wir müssen kognitive Agenten schaffen deren Verhalten wir verstehen, um Agenten zu erhalten die Informationstypen wie Menschen nutzen und repräsentieren.

Um das Problemlösen auf einer Granularität von Informationstypen zu untersuchen, werden Aufgaben benötigt, die auf verschiedene Informationstypen reagieren. In dieser Arbeit definiere ich die Anforderungen an solche experimentellen Paradigmen zur Untersuchung der räumlichen Informationsverwendung beim Problemlösen. Ich

habe die folgenden drei experimentellen Designs entwickelt, in denen ich den Ansatz der Informationsgranularität auf verschiedene Fragen zum menschlichen Problemlösen angewendet habe.

1. Eine Analogie-Aufgabe, um zu testen, ob Menschen in der Lage sind, ihre räumlichen Fähigkeiten auf neue, nicht-räumliche Domänen anzuwenden.

2. Eine Tic-Tac-Toe-Isomorph-Aufgabe, um festzustellen, ob Problemlösende räumliche Informationen zur Unterstützung in eine nicht-räumlichen Aufgabe suchen. Und wenn ja, würden räumliche Informationen, die für die Aufgabe schädlich sind, die Leistung behindern oder ignoriert werden?

3. Eine Kartensortieraufgabe, um zu überprüfen, ob Problemlösende räumliche Informationstypen gegenüber nicht-räumlichen Typen bevorzugen, wenn beide gleichwertig gute Lösungsansätze sind.

Zusätzlich baue ich ein Computermodell, um zu veranschaulichen, wie meine experimentellen Paradigmen zum Testen und Entwickeln von kognitiven Agenten verwendet werden könnten.

In dieser Arbeit zeige ich, dass Menschen in der Lage sind, räumliche Informationen in nicht-räumlichen Aufgaben zu nutzen, falls diese ihnen bei der Lösung einer Aufgabe helfen. Ich zeige weiter, dass dies nicht nur ein Effekt davon ist, dass räumliche Informationen auffälliger sind oder bevorzugt werden. In einem Setting, in dem räumliche und nicht-räumliche Informationstypen in Bezug auf das Lösen eines Problems gleichwertig waren, wurde keine von beiden konsequent bevorzugt. Problemlösende ignorierten bereitwillig räumliche Informationen, wenn sie in Bezug auf die Struktur des Problems in Konflikt standen. Ich liefere ein Proof of Concept dafür, dass die computergestützte kognitive Modellierung, angewandt auf die Informationstypengranularität, bei der Untersuchung des menschlichen Problemlösens von Nutzen ist. Dieser Ansatz erleichtert den Aufbau von Computer-Systemen, die Informationen so repräsentieren, dass sie mit menschlichen Repräsentationen übereinstimmen. Er kann angewandt werden, um Feedback-Systeme zu erstellen, die Informationen in einer Modalität präsentieren, die Menschen leicht in ihr mentales Modell integrieren können. Dadurch wir deren Verständnis vereinfacht.

Acknowledgments

I am very grateful for the support and opportunities given to me by Christian Freksa. The majority of work for this thesis was done in his research group Cognitive Systems (CoSy). Christian provided me with the opportunity to broaden my horizon after my undergraduate studies. He expected his doctoral students to find their own research questions and through enthusiasm and curiosity encouraged me to find my own way. I am especially thankful for the creative and fruitful exchanges about arising questions, hypotheses or experiments. It pains me greatly that Christian will not be able to see me finish my thesis, he will be sorely missed.

Accordingly, I very much appreciate that Holger Schultheis was willing to take over as supervisor. Bar that, Holger deserves my gratitude for countless times spend brainstorming ideas, scrutinizing experimental paradigms and discussing results. The outcome of which are the experiments that form the basis of this thesis. In times that I could not imagine finishing this dissertation and you helped me see a path.

I also want to credit the whole team of CoSy and the SFB/TR 8 Spatial Cognition for providing such an inspiring environment. It felt like being a part of something bigger. Special thanks to the great support by Grazia Kranz and Alexander Middelstorb. I warmly remember the fun days in the Office with the other doctoral students with whom we navigated the pitfalls of post gradual research. Special thanks to Jakob Suchan for welcoming me in Bremen and the amazing time sharing the office. As well as Immo Colonius for sharing many interests even if our research was quite different.

Within CoSy I was part of the R1-[ImageSpace] project. I want to thank Thomas, Holger, Jan and Ana for introducing me to the fascinating research field of mental representations of space and computational modeling. Not less for the evenings spend tasting whisky, such a diverse and fun team to be a part of.

I want to thank my Parents for encouraging and enabling me to try new things and find my way. You supported me through each section along the way. Even if I did not always succeed, it enabled me to take the next step again with confidence. Thank you for taking a great interest in my research and encouraging a bigger perspective. Or simply supporting me by driving over one hundred kilometers just to see me play.

Unfortunately my father also would not live to see me finish this chapter of my life. I am eternally grateful for your companionship along the way.

Thanks to Basti, Boris and Simon for all the great memories made together that lead me here. And for providing home base in Hannover. And to the Floorballers at Eiche for providing much needed balance.

Finally I am most thankful for the support of my wife Marloes Wienemann-Everlo. I am forever indebted to you for without you I would not have been able to finish this. You sacrificed a lot for us and joined me in Bremen. And I very much appreciate your courage. We married and have a great family together. You are a great mother and make me better. I am exited to begin to return the favor and be able to support you with your next adventure. I will be by you side as you were at mine, thank you.

Contents

List of Figures

List of Tables

1. Introduction

This thesis is about cognition, more specifically how we think about cognition, how to approach studying cognition. Cognition is fundamental for purposeful, goal-oriented behavior. An agent must perceive the environment, interact with it, and store and process knowledge about it. To be successful, none of these skills alone is enough. An agent that encompasses all of these abilities is a *cognitive agent*. An agent and its environment form a *cognitive system*. Various theories have been formulated to describe the interactions within this system. In the past, less has been said about what kind of information we retrieve from the environment for goal-directed behavior. This thesis addresses this issue.

1.1. Motivation

Perception and interaction with an environment are central to the functioning of a cognitive system and should therefore be studied in the context of that environment. For the real world, humans are a great example of a part of a highly capable cognitive system. Humans are very efficient at extracting information from the environment or manipulating the environment to support their functioning.

Introspection provides a readily available method for studying the workings of the human cognitive agent. It allows us to analyze and report on our thought processes. This is not without its limitations but has led to human cognition being heavily used as a role model for artificial cognitive systems. Human behavior in general can be seen as a form of (repeated) problem-solving. Problem-solving uses a (sequence of) step(s) to achieve a goal. Imagine I want to enter my apartment after a visit to the farmer's market. The goal would be to unlock the door with a key. The task is then to remember where I left my keyring, balance the bags of produce on one arm to retrieve the keyring from my pocket with the other arm, identify the correct key, align the key with the keyhole, insert it, and turn it. Moreover, each of these steps can even be considered as a solution to a subprHuman problem-solving has been studied extensively in a variety of disciplines, most extensively in psychology and cognitive science. Most research deals with common errors in problem-solving or the strategy used. Far less attention has been paid to the question of what kind of information people use to solve problems. To solve a problem, an agent must assess the situation and find the appropriate response. Humans as cognitive agents do this through perception, i.e, through information retrieval. It would therefore be very interesting to know what information humans retrieve from the environment to solve a problem. In the example above, if I need to identify the correct key from a keyring, I could use the color of the keycaps, the shape of the key, or the position of the key relative to other keys on the ring. Would I use the same information if I were to select a key from a keyring I am not that familiar with? To have a good understanding of human information use, it is important to understand what kind of information people use when performing tasks.

For artificial cognitive agents, this knowledge would be even more beneficial. Artificial cognitive agents, for example, are used as models of aspects of human cognition. These models provide insight not only into what effects exist in problem-solving, but also into why they may occur, and therefore form a very important tool in understanding the human mind. One caveat that exists with this approach is that the information these models use must be coded into the model a priori. In the environment, there are an almost infinite amount of information sources, and choosing the right one for the model becomes crucial for creating a meaningful model. So understanding what information people use in problem-solving would help in these areas to create much more appropriate models.

The artificial cognitive agent could also be designed not as a model, but as a system that a human would interact with. In this case, it would still be very advantageous if the information that the agent communicates with the user is of the type that the user can be expected to have and is looking for to solve the given task. This would facilitate the use of such a system by eliminating unnecessary information processing. This could especially be helpful in domains with time-critical decisions in information-rich environments such as aviation.

Advances in artificial intelligence have revealed capabilities in machines long thought possible only for humans. These novel approaches promise to outperform human performance on many tasks within my lifespan. However, these systems rely on training neural networks, which are inherently black boxes and do not allow us to examine exactly what information is used to solve a problem. Computer-based analyses still offer this possibility and will therefore retain their justification, especially when these systems have to work alongside human operators. Black box systems for decision-making also cause a whole range of moral and ethical problems due to inherent but unclear biases in information processing.

Of particular importance within human cognition is *spatial cognition*, that is, the acquisition and use of spatial information from the environment. Spatial cognition is crucial for essential abilities, from motor coordination to wayfinding. Spatial information is ubiquitous in the environment, and in conjunction with its importance to human functioning, humans have developed very efficient processes for using spatial information; so much so that we even use spatial metaphors to communicate or reason about abstract concepts, such as the number line.

Starting from the human cognitive system and the apparent special role of spatial information within this system, many interesting questions arise. In my thesis I have concentrated on these:

- Is it possible to isolate different types of information and thus study their use in problem-solving?
- Will people use spatial information to reason in a novel non-spatial domain? Is the spatial information detectable in the behavior in the non-spatial domain?
- Do people seek spatial information to solve tasks that do not require spatial information? Can this also have a detrimental effect if the spatial structure is misleading in relation to the problem structure?
- If people prefer spatial information to other types of information for solving tasks, is it simply because spatial information is more salient?
- Would a problem-solving approach based solely on the use of competing different types of information be feasible, and how would this compare to a human problem solver?

1.2. Method

Despite its ubiquitous use, the term information is not as clearly defined as one might think. It is used in different disciplines with different meanings. In order to study what information people use in problem-solving, I first had to create a theoretical framework to formalize the information content and its use. As mentioned earlier, there are an almost infinite amount of different types of information in the environment. These sources of information arise from the inherent structure of the world and its properties. This *world structure* is defined by the inherent properties of the physical environment that allow certain operations and create constraints on interaction with the environment. For example, the spatio-temporal domain but also other physical and chemical properties. This structure is different for agents with different capability prerequisites. For example, a high fall has few consequences for an insect, while it can be fatal for a human. Interactions with the world shape, our world structure, and the resulting structure reflect the difference in interactions between agents.

In our world, an object must have one location in three-dimensional space and cannot be at another location at the same time. This enables, for example, the restriction that you do have to continue searching in other places once you have found an object. Furthermore, if the object is perceived at a later time in a different location, it can be concluded that it must have moved through space to get there. This assumption holds because, in our world structure, teleportation is not a possible operation. The object will also have a color, specific weight, or chemical profile, each adhering to its own valid operations and characteristics.

This world structure consists of *structural elements* that differ from each other in a meaningful way. A distinction arises from which measure one would use to quantify the element. Following this distinction, I define structural elements such as distance, measured in a unit of length like meters, or color, measured in wavelength of reflected light, or duration, measured in seconds. These structural elements are modality-independent in the sense that there are elements that can be perceived by different sensors but retain their inherent properties. Spatial or sonic structural elements, for example, can be perceived by multiple human senses. Structural elements make up the world structure by representing the properties of objects and the relationship among objects. From these structural elements, a viewer can extract information such as color. When considering what information people use, I distinguish information according to the structural element it carries. The *information type* thus classifies the information according to the structural element that the agent it contains obtains from the environment. So, in the above example of the keyring, my decision could, for example, therefore be based on the information types *color* or *location*. Again information types are modality-independent, the information type *distance* can be perceived by sight, hearing, or touch.

Based on this framework of world structure to information type, I designed three experiments to address the research questions posed above. I also created a cognitive model to examine the validity of a problem-solving approach based solely on information types. To set up the experiments, I started from known paradigms in the problem-solving research and used them in a way that allowed me to this novel approach to problem-solving. The advantage of using familiar tasks and paradigms is that tasks are used that have been shown to be useful in human reasoning and problem-solving research.

The first experiment was set up to test whether spatial metaphors are simply learned and therefore used, or whether people are able to reason with spatial information even in novel domains? I have developed a novel experimental paradigm that uses analogies to study reasoning based on one type of information in a different domain. I show that this can be used to test the ability of reasoners to use spatial information types to reason in a non-spatial domain.

I designed the next two experiments to isolate the information types in order to test their influence on problem-solving behavior. A game playing setting was used in these experiments. The second experiment allowed to test whether spatial information is sought by the problem solvers even if it is not needed for the solution and whether this spatial information might not even be detrimental if it is misleading with respect to the problem structure.

The third experiment was set up to test the relative salience of information types and to distinguish between spatial and non-spatial information types. For this task, I also created a cognitive model that solves this task using an approach based on the use of information types. I then compared the performance of this model to that of human participants on the same task.

1.3. Contributions

- I show that the assumption is correct that structural elements of the environment can be isolated such that their effect on problem-solving behavior can be studied.
- I show that people willingly and successfully use spatial information in a non-spatial domain.
- I provide an indication that people seek spatial information from the environment to help solve problems, even if when it is not required.
- I develop a novel experimental paradigm that provides the ability to test cross-domain reasoning based on an information type.
- I create a novel experimental design to test the influence of spatial structure on problem-solving behavior.

- I provide an example of an artificial system that reasons based on the relative importance of information types.

1.4. Structure of the Thesis

I will first provide an overview of the current state-of-the-art theories of human (spatial) cognition and the use of information during problem-solving, and how these are currently being studied. I will then describe in detail the questions I seek to answer in this thesis and how they extend the theories discussed in the state of the art. Next, I describe the process and reasoning behind the design of the experimental task and the cognitive model. Then I will discuss the results I found while conducting these experimental studies. Finally, I evaluate my findings in the light of the theories, and the conclusions I draw from them before concluding by describing the contributions of my work and prospects for future work.

2. State of the art

In this chapter, I describe the starting point for my research. I will first show that problem-solving, and in particular human problem-solving, is an interesting area of research. I will then expand how problem-solving is, at its core, an information-processing and cognitive matter. I will discuss the difference between knowledge inherent in the world and information.

I will introduce the current theories of human cognition and how they relate to problem-solving. For this, I will discuss the computational approach. It will become clear how much importance modern theories attach to input from the environment.

From the knowledge inherent in the environment, spatial information is constantly collected by people. The omnipresence and importance of spatial information for functioning lead to much research in spatial problem-solving. I will further discuss the implications of spatial cognition.

A prominent way to study cognition, in addition to empirical studies, is through cognitive modeling. The availability of computers and the advantages over empirical studies have made this approach a staple of cognitive science research. I discuss the use of cognitive modeling and the various approaches therein.

In the next part, I show that many approaches to studying spatial cognition, whether through cognitive modeling or empirically, operate at higher levels of abstraction in terms of interaction with the world. I will discuss that the level of analysis, as well as the level of abstraction of the information one uses, is not at that of basic spatial invariants in the world. I conclude with Shepard's call to find cognitive invariants of the world, immutable truths that arise from the structure of the world that influence human behavior. It will become clear that most research focuses on a higher level of abstraction than invariant structure.

2.1. Problem-Solving Research

I examine the role of spatial information in problem-solving. The problem in this domain describes a goal to which the path is not yet known. The term problem describes a wide range of tasks. It may mean getting from home to the train station on time, it may mean preparing dinner, solving math homework, but it may also mean developing new tax legislation or writing a story. There is a difference in these examples; the solutions to the earlier problems such as reaching the station on time can be described as a series of steps in an algorithmic manner. The steps may not be known and there could be various steps that could be taken, but these can be easily formalized in the form of such steps. In the case of reaching the train station, this could, for example, be calling a taxi, waiting for the taxi, and then going to the station by taxi. The latter examples, on the other hand, are not so easy to describe in steps. Moreover, the target state for these problems is not quantifiable. The quality of tax legislation cannot be determined objectively. Too much depends on political views and personal interest. Likewise, the question of what is a good story will be answered fundamentally differently by me than by you or any other person. The methodology from management science/operations research was not appropriate for these types of real-world problems. H. A. Simon and Newell (1958) distinguished between well-structured problems and ill-structured problems. Well-structured problems are those that are quantifiable, formalizable, and solvable by an algorithmic routine. The authors claim that the most important everyday problems tend to belong to the poorly structured problem spectrum. They found that the solution was a heuristic-based approach to problem-solving. In doing so, they focused on how humans approach problems rather than trying to explain human behavior in the limited way that technology could solve problems (H. A. Simon & Newell, 1958). The study of human problem-solving has become a fundamental part of what became the cognitive sciences. A mixture of approaches and methodologies is used to describe and explain phenomena. Since the beginnings of empirical psychology (Donders, 1869; Wundt, 1874) and the subsequent behavioristic movement, which studied mainly small-scale effects because of the limitations of the experimental rigor required for their methodology, new disciplines brought novel methodologies to study problem-solving, such as observational protocols, think-aloud protocols (Ericsson & Simon, 1980; Lewis, 1982), and modeling (Sun, 2008).

Problem-solving provides a great research area for real-world problems. Much of human behavior can be described as a form of problem-solving. When one encounters a novel situation, one must adjust one's behavior to achieve a goal. Some of the problems we encountered we solved at a young age early in development, like standing up and walking, others later, like arithmetic. We have learned solutions that allow us to not perceive these actions as problems. Other problems we will never be able to solve due to limited mental capacity and the amount of computation required. For this thesis, I follow Kotovsky and Simon (1990) in defining problem-solving as the use of a *series of steps to achieve a goal state*.

Problem-solving can be seen as a form of information processing and a cognitive task. I will discuss both in the following parts. It can be seen as information processing or information used in the sense that knowledge from the world is used, evaluated, changed, and combined. Problem-solving becomes a cognitive matter as it relies on the input, processing, and response of information by the problem solver.

2.2. Use of information during problem-solving

Our central interaction with the inherent knowledge in the world is to extract information from it. The term *information* is to be differentiated from knowledge or facts, as the knowledge available in the world. Although the term information is frequently used and intensively studied, it is ambiguous. Several approaches to defining information exist, ranging in approach from philosophical (Campos, 2009), interdisciplinary(information science) (Capurro & Hjørland, 2003), and technological perspectives (Shannon, 1948). Adriaans (2012) gives an overview of different qualitative as well as quantitative definitions of information. The essence of information is that it is a transfer of knowledge. Information transports a fact within a signal over a channel to a receiver. Knowledge in the world only becomes information when, for example, it is perceived.

In their influential work, Shannon (1948) formalized information transmission as a probabilistic model within the entropy of a signal with noise. The formalization was intended for telecommunications implementations but was soon adopted by cognitive psychology. The approach of formalizing information within a noisy channel was tantalizingly close to real-world fuzzy perception. Shannon was not concerned with the content of the information but with the underlying mathematical model. For psychological research, the model was adapted also to take into account the content of the information (Wagemans, Feldman, et al., 2012). Other adjustments were made to account for the complexity of natural language (Miller, 2003) and scene interpretation (Hatfield & Epstein, 1985).

In the context of problem-solving and cognitive psychology, the content of information and retrieved knowledge is central. Searle (1980) demonstrated the importance of content by pointing out shortcomings of artificial intelligence systems. They propose the following setup to illustrate that a computer executing a program cannot have intentionality, mental states directed toward or about objects, such as beliefs or intentions.

There is a machine that represents a room with an input slot and an output slot. Inside the machine sits a person. The person has a lookup resource to find corresponding English words to Chinese symbols. The machine is given a slip of paper with Chinese symbols as input, and the person uses the lookup resource to translate the text into English. The person does not need to understand either text to complete the task. To the user, this machine may look as though it is *intelligent* and understands Chinese and English, but Searle argues that formal operations, like looking things up, do not constitute intentionality and therefore cannot be considered intelligent.

According to this view, the system lacks an understanding of the content of the messages it transmits and therefore cannot be considered truly intelligent. To make meaningful inferences about human information use during problem-solving, the content of information should be relevant. If we consider the person in the Chinese room as an experiment who translates without comprehension, that person would not convey information about the content because they lack the processes to break down the facts from the environment into meaningful information. At the level of perception, the shape of the Chinese symbols and the English letters conveys information to the person. The content of the information is on a different level of understanding. The crucial difference is the level of abstraction at which the information is transmitted. Searle might argue that Intentionality requires information at a higher conceptual level, with which I agree.

Real-world problems involve such a wealth of influences and possible sources of information that a complete representation of a problem is impossible for all but trivial problems. Human problem-solvers thus work under the pretext of incomplete information. H. A. Simon (1972) introduced the concept of bounded rationality to describe this. Bounded rationality for human problem solvers means they lack the mental resources to represent all available information. The problems themselves may be intractable, making it impossible to represent the relevant information in full. Simon describes the strategy of problem solvers to deal with bounded rationality *satisficing*. Problem solvers produce a solution that is good enough, if not perfect. Humans are indeed capable of very good performance on well-known computationally intractable problems such as the Traveling Salesman problem (MacGregor & Ormerod, 1996), which in certain cases are close to optimal. Some instances of NP-complete problems are played as recreational puzzles (for an overview, see Kendall et al., 2008). Simon's work was groundbreaking in that it underscored that human problem-solvers do not always act or think completely rationally. In later work, they extend this notion to the use of heuristics as tools in information-rich problem-solving situations, as described above.

The difficulty of a problem may differ by the complexity of the main problem or the size of the problem instance to be solved, but not exclusively. Kotovsky and Simon (1990) studied the differences between problems that were similar in structure but perceived to be very different in difficulty. The puzzle the researchers used to illustrate this is called *Chinese rings*(The puzzle used is an isomorph of *Baguenaudier* itself, also sometimes known as *Chinese rings*). Figure 2.1a shows an example of such a puzzle. The object of the puzzle is to free a metal loop threaded through interlocking rings. The first obstacle to the game is that, unlike games like Tower of Hanoi, it is not clear what constitutes a move. There are many actions to manipulate the puzzle that are not helpful in solving it. This makes this puzzle incredibly difficult. The version shown in Figure 2.1b is a version for consumer products, where the moves are more salient. For their study, the authors created digital isomorphs of this puzzle that were structurally equivalent to the original but clearly communicate what a move is. Only one of the twelve participants solved the puzzle in the allotted time of 90 minutes, another was on the right track at the end of the time. Even with instructions on what a move is, only about half of the participants could solve the puzzle in 60 minutes. The digital isomorphs created for the study were solved by all but one of 26 participants in the allotted 60 minutes. Thus, the difficulty of these problems does not arise from the size of the problem space, as this was the same for both versions and was further supported by the control within the isomorph. The conclusion the authors come to is that the difficulty arises from identifying a move and identifying legal moves. Regarding the isomorphism of the two problems, the authors conclude:

> We defined isomorphism as identicality of the problems' external search space, defined as a set of nodes representing knowledge states, which are linked by paths defined by legal moves. Given this definition, problems whose moves are hard to discover are difficult precisely because the subject's representation early in the problem, is not isomorphic to the representation of the digital problems. As they work on the problem and come to discover what constitutes a move, they achieve a representation that is isomorphic. (Kotovsky & Simon, 1990, p.181)

Another form of difference in representations was discussed by H. A. Simon (1978). The author differentiates isomorphs of representations as *informationally equivalent* and *computationally equivalent*. Informational equivalence between two representations is achieved when all information in one representation is also available or derivable in the other representation. Computational equivalence requires informational equiva-

(a) Original Chinese Rings puzzle. The goal is to free the U-shaped rod.

(b) An isomorphic puzzle. The goal is to remove the red slider from the grey track. *Spin Out* from Binary Arts is now called Think Fun.

Figure 2.1.: Chinese Rings puzzle variants. Both puzzle solutions are isomorphic. There is a difference in the number of “rings”, six in (a) and seven in (b), but the solution is the same. What constitutes a move, though, is less clear in version (a). Each “ring” has a state *on* or *off* the rod(*vertical* versus *horizontal* in (b)). In both versions, only the rightmost “ring” state and the “ring” just to the left of the rightmost *on*/*vertical* “ring” can be changed, with all the “rings” further right being *off*/*horizontal*. In picture (a), Rings One and Two from the right could be changed. In (b), the diagonal and the rightmost knob can be changed.

lence, but in addition, it also requires that any inferences that can be drawn from one representation can also be drawn from the other with similar effort(or, in a looser definition, with proportional effort). A similar effort is not a hard measurement, but rather fuzzy. As an example, Simon used the construction of a geometric figure. There is a

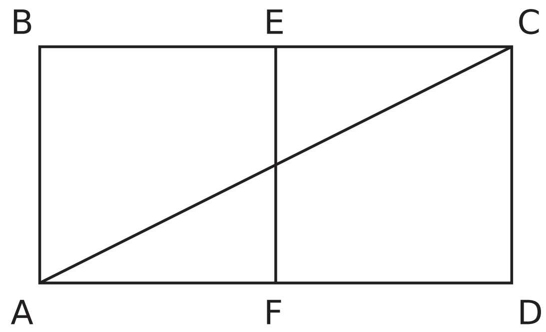

Figure 2.2.: Example of a diagrammatic representation for solving the intersection of AC and EF

Rectangle ABCD with distances BC twice as long as AB. Points E and F are the midpoints along BC and DA, respectively. If there are lines AC and EF, do they intersect? This is one possible representation of the problem, another might be a diagram such as Figure 2.2. Both representations are informational equivalent, but at least for humans, the query is significantly less computationally demanding for the diagrammatic representation. Most people might even translate the first representation into a mental image to solve the query from a pseudo-diagrammatic representation, thus relieving the computational demand.

A prevalent theory of how people solve problems such as those above is the Mental Model Theory (Byrne & Johnson-Laird, 1989; Johnson-Laird, 1989). According to this view, problem solvers create mental representations of a problem on which they operate to reach the goal. A mental model is a simplified representation of (part of) the problem that is not bound to a modality. In fact, the debate about the imagery of mental representations is based on different views about the modality of the representation. The imagery under discussion is that which one experiences when *seeing* something in one's mind's eye. Two competing theories are whether these images are visual (Kosslyn et al., 2006) or not (Pylyshyn, 2003).
A visual mental image could be considered as a modal mental model.

The mental model is an abstraction to ease the cognitive load and allows one to operate mentally with this representation. Abstraction means that less information needs to be stored, so conclusions based on the model more efficient. This abstraction can lead to biases or even fallacies in reasoning (e.g. Tversky, 1981). These effects can provide insight into the structure of the mental model. A well-known example is a high rate of misjudgment as to whether Reno, Nevada, or San Diego, California, is farther west (Stevens & Coupe, 1978). In their study, almost all participants misjudged the direction, thinking San Diego was further west than Reno. The confusion stems from the fact that Nevada is generally east of California, while their shared border is not exactly north-south. It is assumed that the false judgment is an artifact of a hierarchical representation of the environment by the participants, following the reasoning Reno is in Nevada, Nevada is east of California, ergo Reno is east of San Diego (Hirtle & Jonides, 1985). Similar assessment errors have also been shown in north-south assessments between Seattle, USA, and Montreal, Canada (Stevens & Coupe, 1978). Another heuristic that leads to biased mental models of geographic locations is that of alignment. North and South America are aligned together, although they have little overlap in geographic longitude, and the Americas are aligned to Europe and Africa, although Europe is generally farther north than the United States Tversky (1981).

Mental models have been used extensively in the research area of spatial reasoning. The amodality of spatial information and knowledge makes an amodal model very suitable for reasoning with spatial information. Nonetheless, the structure of the spatial mental model may retain some features of what is represented. Sloman (1971) refers to these as spatio-analogical representations. The structure of the presentation corresponds analogously to the original situation. A highway map, for example, is a spatio-analogical representation of the real world because some spatial structures, namely cardinal directions and relative distance, are preserved. In contrast, a list of intersections with coordinates would not be spatio-analogical (Sloman, 1971, 1975). Schultheis et al. (2014) show that a spatio-analogical mental model can explain effects in reasoning about cardinal directions. Indeed, they show how preferences can arise in the mental model. It has been shown that when different mental models are possible to describe a configuration, people generate only one (Cox & Brna, 1995) not unlike satisficing. If this model is deemed a valid representation of the modeled configuration, the reasoner accepts this model. Some model variations are constructed more frequently by problem solvers than others; these are called *preferred mental models* (Jahn et al., 2007; Knauff et al., 1995). Thus, a problem solver obtains information from a problem representation and combines it into a mental model that enables the operations necessary to solve the problem. The modality of such a model need not be the same as the channel of information acquisition (Giudice et al., 2009; Liesefeld & Zimmer, 2013).

2.3. Cognition

As mentioned above, the problem-solving process can also be seen as a cognitive process. In my thesis, I will focus on problem-solving from a cognitive science perspective. The central research questions are considered with the cognition involved in problem-solving. I mainly consider human cognition, but the issues raised in the thesis are also applicable to other species or artificial agents. I use the following definition of cognition: "Cognition is the collection of mental processes and activities used in perceiving, remembering, thinking, and understanding as well as the act of using those processes" (Ashcraft, 2005, p. 11). This definition is broad as the field of cognitive science and cognitive psychology. By definition, cognition deals with perception, mental processes, and the interaction between them. Cognitive research goes beyond simple S-R effects but is interested in what is perceived and how what is perceived is used (e.g. Navon, 1977).

2.3.1. Computational Approach

From the dualistic view of a separation of mind and body and perceiving and acting being different from thinking (Robinson, 2016) arises the basis for *cognitivism*. The hypothesis is that the mind can be viewed as a symbol-manipulating entity that follows explicit rules. This view is present in the mind as a computer metaphor and is still used today.

The origin of this hypothesis lies in the formalization of computation by Turing's (1937) eponymous hypothetical *Turing machines*. This formalization of computation allowed us to argue that mental processes can also be represented as symbolic processing. *Artificial intelligence*(AI; even in the weak case (Searle, 1980)) is based on this principle of cognition by symbol manipulation.

Turing (1950) proposed a test in the form of a game to determine whether a computational system can be considered to behave intelligently. In their famous proposal of the *imitation game* that became known as the *Turing test*, they described a way to assert artificial intelligence. The setup has been proposed as a test for general artificial intelligence, similar to Searle's notion of strong AI. And much like the Chinese room experiment, the setup illustrates that is it is not necessary to create functionally equivalent machines to achieve intelligent behavior. The original premise of the Turing test was that there is an interrogator who can communicate with two agents A and B. One would be another human, the other the artificial system. The goal for the interrogator was to find out which agent was male and which was female. The communication between interrogators and agents had to be done, of course, in such a way that no clue other than the messages and their timing could be used to deduce the answer. An intelligent system would be one that could play this game as an agent and fool humans. The test was interpreted as whether a human interrogator could tell whether the interlocutor was a machine or a human. It has been argued that this is the intended interpretation of the test (Piccinini, 2000). Alternate versions include detection of *bots*, i.e. artificial intelligence non-player characters, in computer games (Hingston, 2009). The performance of artificial systems in imitation game-like settings is still very poor (Shah & Warwick, 2010). Probably the most widespread application of the Turing test, or a derivative of it, turns the premise around, enables systems to recognize human operators by exploiting the errors of machines to mimic human performance in a perceptive task such as checking on the World Wide Web. The technology called CAPTCHA (for Completely Automated Public Turing Test To Tell Computers and Humans Apart) is widely used to prevent bots from entering online forms (von Ahn et al., 2003).

The computational approach similarly argues that such artificial procedures are representative of cognitive processes and as such, cognition can be explained in the language of computational technology. Newell and Simon (1976) proposes in a seminal work that physical symbol systems are sufficient and necessary for intelligent behavior. The proposed system is a classical symbol manipulation system, similar to how a computer acts. Since the inception of cognitive science as a discipline, symbolic reasoning has been hypothesized as the way the mind works (Newell, 1980; R. A. Wilson & Clark, 2009). These theories take the form of input, computation/manipulation, and output. The way the symbol is entered is not relevant since the processing is independent of the input modality.

2.3.2. Dynamic Systems

The view of isolated processing within the mind has since been challenged by proponents of various theories of cognition such as the dynamic systems approach to cognition (Schöner, 2008) and situatedness of action and cognition (e.g. J. S. Brown et al., 1989; Gallagher, 2008). A dynamic system is one that regulates itself based on interaction with its environment. A prominent example often used in cognitive science is a mechanism called the centrifugal governor. This device was popularly used in the nineteenth century for regulating steam engines. The flow of steam from the boiler to the prime mover was controlled by a valve to prevent the engine from overspeeding. This valve must be constantly adjusted to the speed and the steam flow in order to achieve constant speeds. For a human, this was nearly impossible and required constant attention. The solution to this problem was the centrifugal governor, a device that was turned by the steam engine. Two weights were lifted by the centrifugal forces by turning the lever to which these weights were connected at the opening of the steam control valve. So increasing the speed would raise the weights, which in turn would decrease the opening of the valve and decrease steam flow and therefore the speed. With this simple device, the steam engines reached almost constant speeds and had no danger of overspeeding, even without constant attention. This centrifugal governor will only work if all the subsystems work together and are interdependent. The system also does not rely on displays to regulate the speed of the engine. Rather, the governor and its environment form one dynamic system, which should therefore be studied as such and not as parts. The claim of a dynamical systems approach to cognition is that, like the governor cognition, is such a dynamical system in interaction and interdependence with its environment and as such must be studied in its entirety (van Gelder, 1995, 1998). This approach challenges the inherent state-based reasoning and cognitive theory of Newell and Simon with the computationally less expensive dynamic system explanation:

> The heart of the dynamical approach can be succinctly expressed in the form of a very broad empirical hypothesis about the nature of cognition. For decades, the philosophy of cognitive science has been dominated by the computational hypothesis, that cognitive systems are a special kind of computer. This hypothesis has been articulated in a number of ways, but perhaps the most famous statement is Newell and Simon's *Physical Symbol System Hypothesis*, the claim that physical symbol systems (computers) are necessary and sufficient for intelligent behavior (Newell and Simon, 1976). According to this hypothesis, natural cognitive systems are intelligent by virtue of being physical symbol systems of the right kind. At this same level of generality, dynamicists can be seen as embracing the Dynamical Hypothesis: Natural cognitive systems are dynamical systems, and are best understood from the perspective of dynamics. Like its computational counterpart, the *Dynamical Hypothesis* forms a general framework within which detailed theories of particular aspects of cognition can be constructed. It can be empirically vindicated or refuted, but not by direct tests. We will only know if the Dynamical Hypothesis is true if, in the long run, the best theories of cognitive processes are expressed in dynamical terms. (van Gelder &

Port, 1995, pp. 4–5)

In contrast to the "mind as a computer" approach, cognition in dynamic systems cognition is emergent from non-cognitive subprocesses that act interdependently. Using mathematical methodology to analyze dynamical systems, emergent cognition and behavior are explained (McClelland et al., 2010; Schöner, 2008). Dynamical systems have been described from the perspective of classical symbolic cognitive science as bottom-up versus top-down approaches (McClelland et al., 2010). Dynamical systems approaches have been used to investigate low-level course correction actions in flies (Schöner, 2008), motor development in children (Thelen et al., 2001; van Geert, 1998), but also higher-level cognition such as language acquisition (de Bot et al., 2007). Criticism of this approach has been voiced on the grounds that the dynamical systems methodology merely allows a formalization of dynamic interaction and thus does more to describe behavior than to explain it (Eliasmith, 1996; Leeuwen, 2005). Leeuwen puts it as follows.

> (. . .)there might be some explanatory potential there, but the main thrust of the project appears to be description – a quantified account of certain features of successive stages in an in itself fluid developmental process, with the express intent of enabling comparison between subjects. (Leeuwen, 2005, p. 279)

2.3.3. Situated cognition

Proponents of *situated cognition* acknowledge the importance of the context in which cognition occurs but do not adhere to the emerging concept of the dynamic systems approach. Situated cognition can therefore be classified between the classical approach and the dynamical systems approach. Central to situated cognition is "that cognitive activities should be understood primarily as interactions between agents and physical systems and with other people" (Greeno & Moore, 1993, p. 49)

If we consider again the example of trying to reach the train station, the cognitive processes are involved in every step of the problem-solving process. Each step depends on the environment, the situation you are in at that moment. If the road curves we cannot, you cannot just keep going straight, or a one-way street eases the decision problem at an intersection.

We filter much of the constant sensory input to reduce noise and make decisions based on important facts from the environment. Nevertheless, there is a lot of input that is perceived unconsciously and yet influences mental representations and processing. This was novel compared to the stimulus-input-computation-output paradigm of the computational approach. Several effects have been found in the mid-twentieth century showing that non-relevant information from the situation can influence the interpretation of a stimulus. Bruner and Minturn (1955) showed that the same glyph can be interpreted as either an uppercase B or a 13, depending on whether it is accompanied by letters or numbers. An illustration of this effect can be seen in Figure 2.3a. Selfridge (1955) illustrated the limitations that computer vision systems experience when confronted with identical glyphs in a different context with the second prominent example of this context effect for glyph perception (Figure 2.3b. In these cases, what is perceived

(a) (b)

Figure 2.3.: The interpretation of the glyphs is context-dependent. Example illustrations of the stimuli designed by Bruner and Minturn (1955), Selfridge (1955)

is unconsciously enriched with parts of the scene. Other effects of this enrichment can be seen in the Stroop interference effect (Stroop, 1935), in which the meaning of words interferes with naming the color of ink in which a word is printed if it is a color name. For example, naming the ink color of the word *blue* printed in red would be more difficult (take longer and be more error-prone) than naming the ink color of the word *tree* printed in red. The meaning of the word is not relevant to the task, but the automaticity of reading causes the information conveyed by the word to be grasped as well. In the Stroop task example, the enrichment occurs between visual perception and a cognitively higher process of word recognition, which is also visual. Another famous effect of perceptual enrichment is the McGurk effect. Here the enrichment is intermodal between visual and auditory input. McGurk and MacDonald (1976) showed that syllable comprehension depends on the (combination of) modalities in which they were presented. In the study, researchers presented participants with audiovisual recordings of an actor's face saying syllables like *ba-ba*, *da-da*, or *ga-ga*. The researchers then mixed the visual component of one recording with the auditory component of another. Participants would correctly identify these components if presented in isolation. However, when the video of *ga-ga* was combined with the audio of *ba-ba*, the majority of participants perceived *da-da*. Here, the combination of multiple modalities in dissonance produced a different experience than the isolated stimuli. In word recognition, the effect of context was studied by McClelland and Rumelhart (1981). The authors created a model that simulates the effect of context on the activation of a letter during word recognition. Their work involved letters that were not fully visible or incompletely visually analyzed. The authors show that the meaning of the letter corresponding to the activation had to be inferred by the context of the other known letters. They give the example of the word *WORK*, where the last letter is partially obscured so that the last letter can be a K or an R. From the context of the other letters, the decision would clearly lean toward the K.

The human perception or scene integration appears to be holistic, in the sense that objects are perceived as a whole, and visual input is contextually grouped. The Gestalt movement provided several principles by which this is done (e.g. Rock & Palmer, 1990). This implies that knowledge is not gathered piece by piece but is grouped similarly to visual information according to Gestalt principles. For example, compare reading comprehension. As an early viral message shows, humans are pretty good at understanding text even if the letters in the words are jumbled, as long as the visual structure of the word and its initial and final letters stay in place (Rayner et al., 2006). Thus, the word meaning is not perceived by sequentially scanning letter per letter, but by perceiving the whole word. Similarly, auditory perception receives its input as a whole, and segmentation appears to be semantically driven rather than sound stimulus-driven. Sound features that indicate natural segmentation points in the speech stream (i.e. pauses) do not correspond to individual phonemes of speech or even to words of a sentence. Nevertheless, humans can easily distinguish individual words in a sentence (Liberman et al., 1967). Automated segmentation of speech streams is not an easy task for AI, even if it seems effortless for humans (Goldman, 2011).

The effects described above show that perception does not mean retrieval of isolated facts from the environment. It is crucial to consider the situation from which the facts have come. Proponents of situated cognition use these findings to argue that cognitive systems are not limited to the cognitive agent, but rather consist of the agent plus its environment (Clancey, 2008). Barsalou (2008, 2009) argues that much of cognition is based on simulation and that this simulation is typically situated rather than isolated. For example, if you imagine a bicycle, you would not imagine it abstractly, but in a relevant situation. The situated approach to cognition has also been used to explain phenomena in learning (e.g. J. R. Anderson et al., 1996; J. S. Brown et al., 1989) and the effects of social situations (e.g. Smith & Semin, 2007).

2.3.4. Embodied cognition

Within the framework of situated cognition, a sub theory emerged that deals with the role of the body in situated cognition, *embodied cognition*. The premise of embodied cognition is that the only evolutionary reason for a nervous system is to enable action (Wolpert, 2011). Cognition then is the sensation that the body is acting or simulating an action. The discovery of *mirror neurons* showed that thinking about or observing an action triggers a similar physiological response as performing that action (Rizzolatti & Craighero, 2004). Similar effects have been found in speech perception (Hauk et al., 2004) and interpretation (Glenberg & Kaschak, 2002). Embodied cognition rejects Cartesian dualism but grounds the mind entirely in bodily functions. R. A. Wilson and Foglia (2015) define the embodiment thesis as follows.

> **Embodiment Thesis:** Many features of cognition are embodied in that they are deeply dependent upon characteristics of the physical body of an agent, such that the agent's beyond-the-brain body plays a significant causal role, or a physically constitutive role, in that agent's cognitive processing. (R. A. Wilson & Foglia, 2015, Section 3)

Human behavior shows signs of embodiment, for example, in the tendency to gesture when speaking (Alibali, 2005; Tversky et al., 2009; Wesp et al., 2001), or alpine skiers and lugers who mentally rehearse the course by moving in accordance with the curves (Clarey, 2014, February 22). Similarly to mirror neurons, the type of imagery that is used – visual versus kinesthetic – is represented in the same respective brain regions that are active during visual perception or movement (Guillot et al., 2009). This further supports the link between mental processes associated with their physiological counterparts. For an overview of the use of imagery in sports, see Martin et al. (1999).

Embodiedness of perception and scene interpretation has a strong basis in the Theory of Affordances (Gibson, 1979/1986). In this seminal work, Gibson proposed that the way an organism perceives a scene is fundamentally determined by the way that an organism can interact with the scene. Thus, a flat surface may be perceived as a possible resting place because it has been learned that this is an ideal place to sit. A normal chair could then be perceived as a seat because it allows a person to sit. For an elephant, the same chair might not be suitable for sitting, as its body requires different features for the same affordance. Affordances can determine how we interact with our environment. If trying to open a container with a square lid, you will not try to unscrew it because the shape does not afford to turn. Early strong evidence for the influence of our body on the way we perceive the environment was given by Warren (1984). In the experiment, the researchers asked participants if they thought stair steps of a certain height were climbable. What the researchers found was that the height of the stair steps was strongly correlated with the leg length of the participants. This shows that affordances and scene interpretation depend on the skills of the observers. If we consider reaching the top of the stairs as a problem to be solved, these findings further show that problem-solving behavior is in turn influenced by our capabilities.

Reaching the top of the stairs may seem like a very simple problem. However, more complex actions have also been shown to be solvable using embodied cues that would otherwise be computationally expensive. One such question is how do baseball players know where to run to catch a fly ball, i.e. a ball hit high into the playing field? An experienced outfielder will begin to move towards the ball's landing spot almost immediately after a fly ball is hit. One theory of how they know where to run is that they do internal calculations in the parabolic trajectory of the ball and calculate the landing position from that. This would have to be unconscious since players do not report solving the task this way. Again, this is computationally expensive and would be unreliable given environmental factors such as wind and ball spin. Also, this strategy would allow players to reach the touchdown spot before the ball is there, but players seem to reach the ball on the run. McBeath et al. (1995) showed that players solve this problem using a much simpler embodied approach. Instead of calculating the flight path, the players maintain a *linear optical trajectory*. That is, the player positions himself so that the ball travels on a linear trajectory through his field of view in relation to the horizon and home plate. In this way, the player arrives at the touchdown location at the same time as the ball, without the need for complex calculations. This strategy reduces the calculation to a retinal image interpretation. Similarly, it has been argued that retinal image changes(τ) are responsible for the time to impact control and subsequent wing retraction in diving birds (Lee & Reddish, 1981). For *time to impact* reactions, it has been argued that τ must be augmented by other information in tasks such as catching small or irregularly shaped objects (Tresilian, 1999).

These accounts of embodied solutions for real-time, highly dynamic reactions support the benefits of embodied solutions in situations where time is of the essence.

Evidence for embodied processes has also been found in processes of higher cognition, such as language comprehension. Similar to the context effects in perception described above, Glenberg and Kaschak (2002) showed that responses that were consistent with a described action were faster than those that contradicted it. To test this, the authors had participants judge the sensibleness of the sentence by pressing a button press away from or toward the participant. The sentences that were assessed had an implicit movement that was consistent with or contradicted the response direction. The results show a strong interaction between response direction and implied motion in the sentences. The authors take this as strong evidence for the embeddedness of language. Havas and Matheson (2013) further propose that neural systems for experiencing emotions are also relevant for understanding emotion in language. Restriction of facial expressions by BotulinumToxin-A(BTX) has been shown to inhibit emotion recognition (Havas et al., 2010). This happens selectively for emotions that use the subdued muscles for their respective facial responses. In the case of the referenced study, BTX treatment inhibited the recognition of sadness and anger while it did not affect the recognition of joy since the muscles that were subdued were those used in frowning. The authors argue that this underlines the functional role of peripheral motor systems in cognition. Reduced muscle activity caused by high levels of exercise exhaustion has also been shown to reduce the perception of affective state (Brunyé et al., 2013).

M. Wilson (2002) does not summarize Embodied Cognition as one theory but identifies six claims that at least some make for the theory. Namely (1) Cognition is situated, (2) Cognition under time pressured, (3) We off-load cognitive work to the environment, (4) The environment is part of the cognitive system, (5) Cognition is for action, and (6) Offline cognition is body-based.

(1) Embodied cognition emerged from situated cognition and thus accepts that cognition is situated. (2) In addition to situatedness, and this is also often part of the situated cognition theory, cognition works *online*. Any theory of cognition must account for continuous-time and temporal constraints. (3) People use their environment to reduce cognitive load. For example, we count off on our fingers, take notes, or retrace our steps when we are looking for something. (4) Human cognition is so intertwined with the environment that this environment can become the entire cognitive system. This is similar to the arguments used in the dynamic systems theory of cognition discussed above. M. Wilson (2002) doubts the validity of this claim. (5) A nervous system is needed only for coordinated actions. A popular example is that of a sea squirt that digests its nervous system as soon as it becomes stationary with maturity (Wolpert, 2011). (6) Even in absence of the stimuli from the environment, cognition is based on processes that have evolved for perception and action.

2.3.5. Real-world interaction and cognition

What is apparent from these more modern theories of cognition and the effects of situatedness is a great emphasis on input from the environment and the role that this interaction plays in cognition. Of the dynamic systems, situated, and embodied cognition are all based on the fact that information from the environment or learned interactions with the environment should be considered when studying cognition. One type of information from the environment that is important for functioning is spatial information, be it for wayfinding or purposeful movement. This fundamental role of spatial information for (human) functioning is the reason why I will focus on the use of spatial information in this thesis. In the next part, I will describe how spatial information has been studied in cognitive science.

2.4. Spatial cognition

A central insight from situated and embodied accounts for cognition is the role of space. If we accept that all cognition is situated, the role of space in cognition becomes the core of the study of cognition. We live in a three-dimensional world, and all interactions with the world necessarily have a spatial aspect. Spatial cognition is thus a fundamental part of goal-directed action. It is essential for foragers and hunters to find food in order to survive and essential for most animate organisms to reproduce. Over time, it is also omnipresent, and one cannot detach oneself from it. Hart and Moore define spatial cognition as "the knowledge and internal or cognitive representation of the structure, entities, and relations of space; in other words, the internalized reflection and reconstruction of space in thought" (Hart & Moore, 1973, p. 248)

Spatial cognition can take the form of navigational tasks such as wayfinding (Allen, 1999; Hölscher et al., 2006), route planning/tracking (Allen, 2000; Brunyé et al., 2010; Conroy Dalton, 2003), self-localization (Levine et al., 1984), and the construction of a cognitive map (Hirtle & Jonides, 1985; Kuipers, 1983; Tolman, 1948). However, hand-eye coordination is also a spatial task and thus many manual tasks rely on spatial cognition. Imagine you are mailing a letter, for example. The orientation of the letter to put it in the mailbox must be correct. To do this, the environment must be perceived, and the hand positioning adjusted accordingly. Such tasks seem trivial to humans, but work in robotics has shown that this is no easy feat. While computers can compute and play chess at the level of superhumans, the manipulation of the environment is even less fluid in artificial agents (e.g. Kunze & Beetz, 2017).

The human brain structure further supports the importance of spatial knowledge. Visual processing streams clearly separate into so-called *what* and *where* pathways after the primary visual cortex, where dedicated brain areas primarily process spatial information (where) in the dorsal stream and object information (what) in the ventral stream (F. A. W. Wilson et al., 1993).

Performance on navigation tasks has been shown to be an accurate and efficient predictor of Alzheimer's disease in aging participants (Coughlan et al., 2019). This underlines once again the fundamental importance of spatial cognition for humans.

2.4.1. Embodied spatial cognition

The ubiquity of space and its importance for interaction with the world would imply embodied effects of spatial cognition. The assessment of the scalability of steps described above is an example of an embodiment of interpreting the spatial properties of one's environment. Although some argue that there is no clear evidence that cognition influences perception per se, but that interpretation of what is perceived is influenced (Firestone & Scholl, 2016). There is also evidence of embodied effects in other areas of spatial cognition. A gravitational bias has been described in route planning (Brunyé et al., 2012; Brunyé et al., 2010). Perspective-taking distinguishes which *frame of reference*(FoR) one uses to interpret a scene, with a preference for a body-centered FoR noted. Levinson (1996) distinguishes three FoRs: The *relative* FoR that depend on the perspective of the perceiver, such as "he's to the left of the table", the *intrinsic* FoR, which uses features of an object such as "in front of the car", and the *absolute* FoR which uses a global reference system such as cardinal directions, "London is north of Paris". Note that the usage is not always clear. If, for example, someone says "in front of the car", he could use either the relative or intrinsic FoR; in the first case, the described object is between the car and the observer. These FoRs were classified into two categories, namely *egocentric* and *allocentric* (Tversky & Hard, 2009). Egocentric FoRs are those that use relative terms such as left and right; these include Levinson's relative and intrinsic FoRs. An allocentric FoR uses a global reference system and absolute terms such as cardinal directions and thus corresponds to Levinson's absolute FoR. A strong preference for an egocentric frame of reference has been found to support an embodied perspective-taking process (Kessler & Thomson, 2010; Tversky & Hard, 2009). Notably, there are exceptions in some cultures that favor allocentric frames of reference over egocentric ones. It turns out that an embodied approach to perspective taking is not the universal approach for humans but rather is culturally moderated (Levinson, 1996, 1998; Majid et al., 2004).

Tversky and their colleagues described three different spaces when considering spatial cognition. These are functionally related to the human body and its own interaction with the environment (Tversky, 2009; Tversky et al., 1999). They distinguish between the space of the body, the space around the body, and the navigational space. The space of the body acknowledges the special role that bodies play in perception. It describes the space of one's own body through proprioception or the body of another. It deals with the relative position of the body parts. The space around the body is the one that can be directly perceived and interacted with. This space is subject to asymmetries and peculiarities of the human body and therefore supports the idea of an embodied spatial cognition. This space can also be changed by the use of tools. Neuroscience evidence shows that tools that increase the reach change the interpretation of objects that would be out of reach to be in a range around the body (Bonifazi et al., 2007). The navigation space deals with the part of the space that cannot be perceived from one's own position but has to be explored by traveling and composed of several scenes. In this thesis, I am primarily concerned with the space around the body, that is, the small-scale space that can be manipulated but is not perceptually motor-derived.

2.4.2. Spatial problem-solving

The use of spatial information is very important for everyday functioning and goal-directed behavior. In addition to the navigation tasks mentioned above, other research interests include the mental representation of space (e.g. Noordzij & Postma, 2005; Noordzij et al., 2006; Tversky, 1981), mental rotation (e.g. Hyun & Luck, 2007; Liesefeld & Zimmer, 2013; Shepard & Metzler, 1971), and mental imagery (e.g. Sima et al., 2010). In the mental imagery debate, the debate about the nature of the structure of mental images, it is generally agreed that the basis for an image is the spatial structure of that image. This can be simplified or distorted but provides the framework on which the image is created. Opposing theories exist about the representation of this spatial frame and the generation of the mental image, most prominently the quasi-pictorial theory (Kosslyn et al., 2006) and the descriptive theory (Pylyshyn, 2002).

Spatial information is also used as a tool to represent and think about abstract phenomena such as time (Casasanto & Boroditsky, 2008; Gentner, 2001) and math (Izard & Dehaene, 2008; Moeller et al., 2012; Stoianov et al., 2008). Because of this use of spatial information or metaphors for abstract concepts, a high spatial ability is associated with higher achievement in STEM subjects (Science, Technology, Engineering, and

Math) (for an overview, see Wai et al., 2009). Spatial problems are part of most common intelligence tests (e.g. Wechsler Adult Intelligence Scale-IV, Woodcock-Johnson 4, Stanford-Binet 5, Wechsler Intelligence Scale for Children).

In problem-solving, even the non-abstract tasks are often at least partially spatial in nature, either because they are physical problems that require manipulation of the environment, or because their representation is a physical image, making them at least partially spatial. In diagrammatic reasoning, the problem is solved spatially via the inspection of the problem representation (Funt, 1980). In this case, one relies on low-level comparison rather than computation.

Games provide a form of problem-solving that is structured and therefore suitable for research. Many games include spatial elements as part of their rules. Some of the earliest recorded games are games with spatial *movement toward a goal*. From the ancient Indian *Pachisi* (W. N. Brown, 1964), the Egyptian *Senet* (Sebbane, 2001), the *Royal Game of Ur* popular in the Middle East (Finkel, 2007), and the *Mancala* games played in variations around the world (de Voogt, 2003; Gobet et al., 2004).

Examples of explicitly spatial tasks in problem-solving research include mental rotation tasks, *Tower of Hanoi*(ToH) (Handley et al., 2002; Kotovsky et al., 1985) and its derivate *Tower of London*(ToL) (W. K. Berg & Byrd, 2002; Shallice, 1982), some Raven's progressive matrices, Tangrams (Goodman et al., 2016; Olkun, 2003; Pande & Chandrasekharan, 2014), route finding/planning tasks, mazes, Fifteen puzzle or sliding block puzzles like *Rush Hour*[1] (Archer, 1999; Steffenhagen et al., 2014), Go (Reitman, 1976), Checkers (Epstein et al., 1998; Schaeffer et al., 1995), and Chess (Charness, 1992; Chase & Simon, 1973). The latter games of Go, checkers, and chess have long been a strong topic in AI research to build machines that would be able to beat humans. The ability to play these games has been associated with what makes true intelligent machines. These games are games of total information, all information is visible to all players at all times. Also, the games are *impartial*, which means that the only difference between the two players is who starts. They are difficult because the problem space of these games is huge. Still, people were able to compete at a high level. Computer science had to look for ways other than brute force. The computationally simplest of the games is checkers with roughly 5e20 possible positions solved computationally (Schaeffer et al., 2007). The other two are unlikely to be resolved. Chess and Go are too big to be solved with about e46 and e172 positions, respectively. The search tree for all possible move sequences from a position is much higher still (van den Herik et al., 2002). Problem-solving in these games cannot be solved by perfect knowledge,

[1] ThinkFun, Inc.

and therefore human-like methods are used instead, advocating the power of human problem-solving skills. When IBM Deep Blue beat Kasparov in 1997, it was still using a brute force approach to openings and possible positions and had to be adapted from the version that lost in the year before (Campbell et al., 2002). New versions of chess engines have become even more powerful, and current engines have an estimated Elo rating much higher than that of top human players (Sandin, 2020).

In Go, the first victory of a machine over a high-ranking human on a full board without a handicap occurred recently. In Go, it is even difficult to assess who is ahead, and top players are often unable to explain their strategy, relying more on their intuition. A conventional AI approach would not be able to do this given the size of the problem space. *AlphaGo* won against two top-tier human Go players in 2015 and 2016 by combining neural networks with heuristic search methods (Silver et al., 2016).

The AI systems necessary to achieve superhuman performance have had to learn from human behavior because the problem spaces are too large to brute force a solution and rely solely on fast computing power. The ability to recognize patterns and use heuristics to efficiently utilize computational resources has made it possible for systems to operate at current levels. In AlphaGo, performance is also based on patterns learned by simulating millions of games. Due to the nature of these games, these patterns are spatial in nature. And given the constant exposure to spatial information and the ability of humans to deal with spatial information, these performances demonstrate the strength of spatial cognition in problem-solving.

Recent developments in machine learning, particularly deep neural net reinforcement learning, have enabled AI systems to challenge humans in real-time competitive games as well. These systems rely on the neural networks being extensively trained beforehand. Competitive (multiplayer) computer games like *StarCraft II* or *DOTA 2* require quick reactions to real-time (spatial) situations as well as rather long-term planning of the strategy. In quick succession, two AI teams have had outstanding success against human players in these games. *Deepmind*'s *AlphaStar* (Jaderberg et al., 2018) beat a top *StarCraft II* player(Vinyals et al., 2019). *OpenAI Five* beat the reigning world champion team in *DOTA 2*, a five-on-five game that requires extensive tactical cooperation to succeed at a high level.

Spatially invariant knowledge and other properties from the physical world can be used to facilitate processing Kirsh (1995). Think about finding the shortest route to a destination on a map. If we simplify the map into an undirected graph, this problem can be solved computationally in polynomial time using Dijkstra's algorithm. There are even better-customized algorithms for navigation. Nevertheless, the search space can become very large for people to solve this task. People will use heuristics to reduce this search space Brunyé et al. (2012), Brunyé et al. (2010), Turner (2009), Yang and Schwaninger (2011) By eliminating alternatives, the search space of alternatives is reduced, and the optimal path is found more easily. Freksa shows that this problem becomes trivial when real-world boundary conditions are applied. One could represent the graph by replacing the edges with yarn of appropriate lengths and nodes with knots. In the physical world, if one were to hold the start and finish nodes/knots and pull them apart, the physical constraints of the environment would cause the shortest path to becoming visible by tightening, while all longer segments would sag. Freksa uses this as an example of *strong spatial cognition*, all computation is replaced by manipulation of the world and reduced to sensory *lookup* (Dewdney, 1988; Freksa, 2015). Similarly, if you would want to find the longest spaghetti in a package, you could go and measure each one and then look up the highest value. Alternatively, you could use world constraints to push the bundle of spaghetti onto a flat surface and then see which one comes the highest, again eliminating the need for memory and computation (Dewdney, 1988; Kirsh, 1995). Dewdney dubbed these mechanisms *analogical gadgets*.

2.5. Cognitive Models of Spatial Information Use

Empirical psychology was the most common way to study human cognition in the twentieth century. Experimental setups are easy to perform, and the statistical evaluations allow for comparable and objective results. This empirical and experimental methodology is very good for asserting that a particular behavior exists, but it is limited in explaining why the behavior occurs Kuipers (1983). With the advent of computers and their affordability, a new way of studying psychological phenomena emerged. Rooted in the way models were used in physics and other sciences, the cognitive model became an addition to the arsenal of cognitive scientists.

A cognitive model is an abstraction of mental processes created to describe and explain cognitive behavior. For example, like a mathematical model of gravitation is used to explain the observed behavior of falling objects or planetary movement, a cognitive model attempts to provide insight into complex phenomena. A cognitive model can be viewed as a theory about cognition that can then be empirically tested (Sun, 2008). One advantage of a modeling approach to gain knowledge about a new domain is its higher expressiveness compared to empirical approaches Kuipers (1983). Empirical approaches can only describe the current state, whereas a computational model can provide predictions of how the modeled system would behave in novel situations that have not yet been observed. These can be tested and, if proven accurate, provide a much more expressive result. Furthermore, the abstraction of the model allows testing the influence of controllable and manipulable mechanisms. The model has the ability to explain how effects come about, rather than just describing that they exist. It is this explanatory power that makes them a desirable methodology for understanding cognitive processes.

Cognitive models can be divided into three broad categories: computational, mathematical, and verbal-conceptual. Each has its own advantages and limitations (Sun, 2008). Computer-based approaches use formalized algorithmic descriptions. Mathematical approaches describe relationships between variables with equations. Verbal-conceptual approaches use informal language to describe dependencies. For this thesis, the first two are most relevant. Mathematical approaches can be seen as a subset of computational approaches, but there is a difference in expressive power.

2.5.1. Mathematical cognitive models

Two broad areas of mathematical approaches to modeling cognition are probabilistic models and connectionist models. Probabilistic models have recently been predominantly represented by Bayesian inference approaches (Griffiths et al., 2010; Griffiths et al., 2008). These are based on Bayes' theorem:

$$P(A \mid B) = \frac{P(B \mid A)\, P(A)}{P(B)} \tag{2.1}$$

Bayes' theorem describes how likely it is that a given conclusion is accurate based on prior knowledge about the domain. For example, hearing hooves on the ground would predict horses rather than zebras because our prior knowledge has taught us to expect horses. Because of this fundamental connection to prior knowledge, the Bayes statistic is often used in modeling learning (for an overview, see Griffiths et al., 2008). The use of Bayesian inference as a model of human performance is well-suited to describe learning and interpretation effects in complex situations. It can explain human behavior in a situation with incomplete and fuzzy information.

Connectionist models use neural networks. The recent increase in readily available computing power has made neural networks more affordable for research. A neural network is a bottom-up approach to cognition that is commonly used to model perceptual tasks. A simple neural network works as follows. The networks consist of multiple layers of nodes with connections from nodes in one layer to the next. Each connection has an associated weight. The input layer receives input by activating each node depending on the input, like the sensors of a camera chip by varying light levels. The activation then propagates through the network according to the weights of the connections. The final layer of the network is the output layer, and certain activation patterns there are associated with a particular response. In order to function, the weights of the network must be tuned to a specific function. This is done through learning. The network is presented with a set of stimuli and the expected outcome. Weight adjustment algorithms and repetition allow the neural network to adjust to respond to stimuli as expected. For cognitive modeling purposes, the network would then be run on novel stimuli, and the observed responses compared to an empirical validation study. Neural networks are capable of generating human-like perception, for example, in detection and recognition tasks. The text recognition software is based on neural networks. Part of AlphaGo is also a neural network trained to recognize playing positions. The power of neural networks in performing a cognitive task, especially perceptual and pattern recognition tasks, is undoubtedly strong, but a limitation remains in the explanatory power of the approach. They provide means for artificial intelligence systems to solve a perceptual task but provide little insight into how they do so. Anecdotal illustrations of how this could become a problem exist from early applications of the technology. Allegedly, a neural network was used to detect partially concealed tanks in a wooded area to assist soldiers. The network was trained on pictures with and without tanks and performed well on the practice set. When the model was confronted with novel images, it failed completely. After a while, the developers noticed that the practice pictures with tanks were taken on a sunny day, while the pictures without tanks were taken on a cloudy day. The network had not learned to recognize tanks, but weather conditions (Dreyfus & Dreyfus, 1992). The failure of the system is, of course, something that can be attributed to the way the practice set was constructed, and might not happen in a more vigilant research environment and with more experience with the technology. The core problem of the lack of explanatory power remains. For a cognitive scientist trying to explain a cognitive effect, the black-box nature of a neural network limits its usefulness. Recently, research at Google has attempted to generate the output of a neural network to gain insight into what the network actually learns and does. Researchers were able to produce a visual output of a neural network "hallucinating," which gives some insight into what structures a network responds to but still required suggestive input to produce human-interpretable results (Mordvintsev et al., 2015). Neural networks are used in cognitive science primarily because they allow parallel processing along many "neurons" on the network. Thus, these models are not bound, explicitly or implicitly, by the assumption that cognitive processing is sequential. Neural networks are predominantly used in AI settings due to their lack of explanatory power, rather than for modeling human performance (Griffiths et al., 2010).

2.5.2. Computational cognitive models

The most informative modeling approach for cognitive science is that of the computational model. The relevant models for this thesis are *computational cognitive models*. These are models that describe cognitive processes in terms of computational steps within a computer program. The implementation is less relevant compared to the conceptualization and the interaction of the simulated systems. These models have some difficulty representing or explaining human parallel processing, adaptive or learning behavior, and embodiment, but the explicitness of the processes offers valid advantages (Vernon et al., 2007). Computational systems have not advanced to the point where they can fool human interlocutors into thinking they are human in Turing's imitation game (Shah & Warwick, 2010), but computational dependencies can still provide relevant insights into cognitive processes.

An extended system that encompasses an entire cognitive system, rather than modeling only a part of the interaction, is called a cognitive architecture. For an overview of different types of cognitive architectures, see Vernon et al. (2007). Some well-established computational architectures include ACT-R (J. R. Anderson et al., 2004), Soar (Laird, 2008), and EPIC (Kieras & Meyer, 1997). These have been used to model many aspects of cognition. For example, ACT-R has been used to model various domains of cognition such as cognitive development (e.g. Jones et al., 2000; van Rijn et al., 2003), decision-making (e.g. Marewski & Mehlhorn, 2011; Veksler et al., 2013), reasoning (e.g. Ragni et al., 2010), and problem-solving (e.g. Gunzelmann & Anderson, 2003).

Computational models of spatial cognition also exist and are used e.g. to explain preferred mental model effects (Schultheis & Barkowsky, 2011). The proposed model can conclude about cardinal directions. The authors argue that in terms of cognitive economy, a mental representation of a configuration is made only as detailed as necessary. This was illustrated with a task to infer the cardinal direction. The task was of the following nature. Given the premises “A is west of B” and “C is northeast of A”, answer the question, “What is the direction from B to C?” According to the theory of preferred mental models, people have preferences in the constructed solution. In this task, models representing a prototypical triangle are preferred (Sima et al., 2010). Four

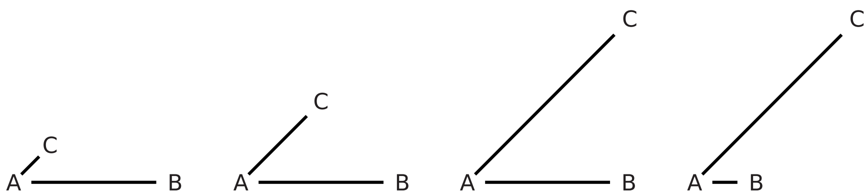

Figure 2.4.: Possible configurations for the premises "A is west of B" and "C is northeast of A". Each leads to a different answer to the question "Where is C as seen from B?"

possible configurations under the given conditions are shown in Figure 2.4. The two middle solutions are predominantly given by people who solve the task. If the task is as follows: "A is west of B" and "C is north of B", the preferred mental model is one similar to the third version in Figure 2.4. Schultheis and Barkowsky (2011) show that these preferences can be explained by an ecological approach to mental representation. The problem solver uses only the levels of detail required to solve the problem. According to the first premise, "A is west of B", only two directions are needed. According to the second premise "C is north of B", the resolution of cardinal directions must be increased to four in order to capture both premises. If the second premise is northwest, the resolution is increased to eight to account for the ordinal directions. The computational approach thus offers an explanation of how a given cognitive effect can be explained by proposed processes.

2.6. Levels of Cognitive Processing

In the sections above, I have discussed the various approaches that can be used to study human problem-solving. In what follows, I will focus on what aspects of a particular process one wishes to study.

Cognition can be analyzed at multiple levels of abstraction. One could study the ability to distinguish colors, which would be at the level of perception of visual input. The analysis could look at the physical processes in the retina. Or you could study the learning progress of a beginner chess player. This analysis could be about remembering openings and applying strategies. Here, the object of study is in a level of patterns and memory. An early and influential hierarchy of analysis for cognitive science is that of Marr (1982). To fully understand an information-processing system, Marr says, you need to understand three levels of the process. The three proposed levels are *computa-*

tional theory, *representation and algorithmics*, and *hardware implementation*. The computational level is the theoretical level that describes the structure of the problem, its goal, and the computations necessary to solve it. The algorithmic level determines the exact steps required to perform the calculations from the level above as well as the representation of the information. At the implementation level, the hardware solution to the problem is described. According to Marr, these levels are largely independent, except for some limitations that one level may exert on another. When describing psychological phenomena, it is important to describe the phenomenon at the right level. Marr provides examples in human vision. To explain afterimages or how colors can be created from three primary colors, one should argue at the level of implementation. Both are artifacts of the retinal anatomy, and understanding these effects requires an understanding of anatomy. While the ambiguity of a Necker cube (an example can be seen in Figure 2.5 Necker, 1832) is not solely due to the perceptual or neurological implementation, rather, the construction of two different 3D interpretations needs to be approached from a higher level.

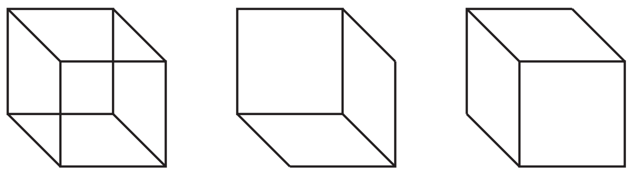

(a) (b) (c)

Figure 2.5.: So-called Necker cubes. The 2D depiction in (a) enables both the 3D interpretations (b) and (c). Human observers can usually switch between interpretations.

Computational cognitive models generally deal with the two higher levels. The neural network approaches also do not claim to model the anatomical processes, even though the concept of a neural network is inspired by the interconnection of neural structures in the brain. Kuipers (1978) builds a model for constructing a cognitive map, a map-like mental representation of an environment, from (route) descriptions. The level of analysis here is one of representation and algorithmics. Reitter and Lebiere (2010) used ACT-R to model path planning that combines visual perception with knowledge integration and retrieval. Similarly, Hiatt et al. (2004) used ACT-R to model perspective taking. ACT-R is not built with implementational accuracy in mind, but rather conceptual and algorithmic analysis. These examples of spatial computational models all focus on the higher levels of analysis, which is also true for the vast majority of models (for an overview, see Madl et al., 2015). Computational modeling generally focuses on explaining the processes that explain certain behavior and effects. The exact biological, physical, or chemical processes that occur in the body to cause the behavior under study are not the subject of the approach.

2.7. Different Levels of Abstraction of Knowledge

Similarly to the level of analysis, the information used may vary in scope and type. A linguist studying the differences between two populations might think about why they pronounce a word differently or why the same word might have a different slant in the populations. In the first case, the differences are auditory or sound-generating. In the latter case, the information is conceptual. There are many different types of knowledge represented in a word, and it is up to the observer to decide which knowledge to extract as relevant information. Consider a milk carton sitting next to a glass on the table. One type of knowledge could be the tensile strengths within the carton to hold the milk or the molecular structure needed to make the glass transparent. Another form of knowledge could be the relative position of the carton to the glass so that it is reflected on its surface. Relevant knowledge differs depending on the method of observation and the intention of inspection. An engineer of material science and an artist might look at the above scene differently and consider the knowledge of different scopes important. The level of abstraction at which the environment is inspected is crucial. If one is only interested in the relative positions of the milk carton and the bowl, one can abstract from the tensile strength and fluid dynamics of the carton and the milk, and the relevant information from the scene is at the level of the objects.

In the above example of tensile strength, the level of abstraction is much lower than in the latter example. Not all knowledge is relevant to a problem solver, and the appropriate level of abstraction is important to limit the amount of knowledge to be retrieved and represented.

Marr and Nishihara (1978) propose an abstraction levels-based approach to representing three-dimensional shapes for recognition by humans. The authors present a process by which low-level visual organization adapts to the required detail. For example, a human figure can be represented at the low level of abstraction as a cylinder with a central axis along the height of the figure. The human figure can be divided into multiples of these basic cylindrical shapes to increase the represented detail. Thus, the figure could consist of one cylinder each for the torso, extremities, and head, and each limb can also be divided into basic shapes, representing digits, for example. There are spatial relations between these parts that allow the modular representation to include more detail in some parts, while the holistic representation has relatively low fidelity. In human perception, the minimum principle has been proposed as a guiding principle that determines the representation needed to describe a scene or interpret a scene. The principle is that when interpreting a scene, the simplest representation is preferred. The unnecessary detail is then abstracted away in this representation. What constitutes simplicity varies and can refer to the form, the information required to describe the object, or the process required to represent the object. This minimum principle in humans is thought to be a mechanism for economizing perception to allow efficient processing (for an overview, see Hatfield & Epstein, 1985). (Lovett & Schultheis, 2014) show that stimulus abstraction can explain response patterns in mental rotation tasks.

In technical solutions for 3D rendering, similar methods are used to increase the efficiency of rendering algorithms. The abstraction level of the representation, in this case, the number of polygons or voxels displayed, is changed depending on the position of an object in a scene. More distant or occluded objects can then be rendered at a lower level of detail to save processing power (e.g. Garland & Heckbert, 1997; Gobbetti & Marton, 2005; Rossignac & Borrel, 1992; Rusinkiewicz & Levoy, 2000)

For problem-solving research, the main level of abstraction is based on objects or the goals of the solver. The intentions and goals of problem solvers are a central interest of problem-solving research due to the discipline's background in economics. For the purposes of this thesis, the motivation of the problem solver is not of interest and will not be considered specifically. For this thesis, the object level is relevant. Questions such as what puzzle pieces the solvers use or what tools they use are common. Sometimes object manipulation is the focus of a task, as in sliding puzzles or Rubik's cubes. Talk-aloud protocols deal heavily with objects, as they can be easily described by the problem solver.

In cognitive psychology research, there is a strong focus on the modality of the knowledge used. The focus on sensory input in cognition lends itself to investigations of the modality with which aspects of the problem are perceived. The computational approach of cognition assumed that knowledge from different sources is integrated into the central processing and then processed symbolically and amodally (Giudice et al., 2009; Liesefeld & Zimmer, 2013).

Psychology research has dealt extensively with the influence of the modality of the information used in problem-solving. The Stroop effect described above is an example of the effect of different modalities. However, further research was also carried out in the field of audiovisual learning. Moreno and Mayer (1999) found an effect of spatial contiguity in multimedia learning that supports the role of spatial information. The research continued to focus on the modality of the information presented rather than the type of information conveyed. An effect of modality on learning has also been documented in behavioral and neural imaging studies (e.g. Ronsse et al., 2010).

2.8. Research Focus on the Use of Information

In the area of information used in problem-solving, the level of abstraction is somewhat broader than the modality. The studies focused on attention, the source of information, and how the information is presented.

The information content researchers are interested in can vary widely. In Shannon's definition of information, the content of the message is irrelevant. The scope of the content of the message can be very different. One difference in message content can be the level of abstraction. Even the same scene perceived by different observers through the same channel can have very different levels of information abstraction. Chase and Simon (1973) discuss the finding that the retention of chess positions observed for five seconds differs between experts and novices. This difference disappears if the position is not an expected possible chess position, but a random one. The authors suggest that the reason experienced players are better at retaining *real* chess positions is that they recall structural information about the position and thus check off piece positions. Beginners do not have this level of interpretation and therefore need to retain many more nuggets of information – position and identity – for each piece. The level of abstraction of the information retrieved by the experienced players is that of the structures of playing positions, while novice players work at the spatial/positional information level. If the position to be maintained were randomized rather than normally reachable, this ability to operate at a higher level of abstraction disappears, and performance falls to that of novice players who never had this ability.

2.8.1. Attention

A common focus within problem-solving research is that of the problem solver's locus of attention. Attention during problem-solving has two facets, that of spatial selection of information sources and that of selection based on modality. Research methods such as eye-tracking and think-aloud protocols focus heavily on which spatial or conceptual areas of the problem the solver is addressing. In think-aloud setups, this is further enriched by insights into the (mental) actions the solver takes to address a problem. What is not addressed here is what type of information (e.g. space, color, or pitch) is used by the solver. Attention research focuses on a higher level of abstraction of information. It is not the type of information that is sought, but the observed object or entity. In Shannon's sense, attention is focused on different channels and sources but is equally unfocused on the content of the message.

Inattentional blindness achieved by selective attention was popularized by Simons and Chabris (1999). Attention may be a process that tunes information retrieval top-down to specific modalities or areas of the perceptual field but again operates at a level of abstraction from object features or larger. In their study, the authors showed participants a video of two teams passing a basketball. One team was dressed in white, the other wore black. Participants were asked to count a team's passes while all players moved chaotically around each other. During the video, a person in a gorilla costume entered the scene. Participants who counted the white team's passes were much less likely to notice this unexpected event than those who followed the black team. The observers were so fixated on the agents in white shirts that the black gorilla never entered their awareness. Inattentional blindness is a key factor in how "magic" works, especially the sleight of hand illusion. The illusionist has the task is of diverting the attention of the audience from the part to be manipulated. When an illusionist brings forth a ball from under a previously empty cup, he does so not by moving faster than the eye can see, but by manipulating what the audience is paying attention to. This is one reason why illusionary magic has been proposed as a scientific research area primarily for cognition (Kuhn et al., 2008). Change blindness (the inability to perceive even significant changes in a visual scene) is closely related to disorders of attentional control. In the flicker paradigm, two versions of an image are presented alternately with a short delay between the two images, alternatively, switch between versions during eye blink. Observers take a long time to notice even large changes due to the lack of attention-directing cues such as movement (Simons & Rensink, 2005). The content of the change seems to be less important than its role in the scene (Rensink et al., 1997).

2.8.2. Information sources

Other research on information use focuses on the information sources and their selection. Human scene interpretation can be described according to *Gestalt* principles. Gestalt psychology proposed a holistic view of human mental processes. For perception, this manifests itself in Gestalt principles that provide for the perceptual organization. The best-known facet of Gestalt psychology is the principles of perceptual grouping. Perceptual organization is most often illustrated with visual stimuli but is not limited to them (Bregman & Campbell, 1971; Szabó et al., 2016). Entities can be, for example, grouped by proximity, similarity in color or shape, similarity in size or orientation, common fate (entities share the same trajectory), symmetry, continuity, or cohesiveness (Wagemans, Elder, et al., 2012). These principles are used to describe phenomena such as that we still perceive a ball partially covered by a leg as a ball and not as two hemispheres. Human perception is set up so that this grouping happens automatically and unconsciously. Therefore, it is often useful to deal with the object-level information sources in a study, as that is a unit into which the perceived scene is divided.

Even in virtual environments, the main interest is the source of information at an object abstraction level, for example, in information search through real and virtual worlds (Fu et al., 2015).

2.8.3. Information presentation

Not all processing of the perceived scene is conscious or intentional. Subconsciously perceived information can still influence the decision-making process. It is important to understand how different presentations of information sources can influence perception and behavior. An example of this was the above isomorphic representation of the Chinese Rings puzzle by Kotovsky and Simon (1990), which showed that isomorphic puzzles differ greatly in difficulty. Informational equivalence is not sufficient to assume that two representation methods are also equally useful as sources of information. The human ability to easily grasp relevant facts of a given representation is based on the processes available to us for processing given representational structures. Early visual processing already encodes orientation and edge detection and allows us to organize visual input. If a given representation does not allow for these mechanisms, the interpretation becomes more computationally expensive, and a possible equivalence is nullified.

In the field of information visualization, the goal is to make the inference you want to highlight from a wealth of data easily accessible. Graphical representations of data relations are one example that can make relations within data more accessible to human observers. The inference example of the rectangle construction in Figure 2.2 is an example of this. Graphical representation reduces the mental inference to visual inspection, which reduces cognitive load because these visual interpretation methods are very efficient. It has been shown that the interpretation and understanding of statistical graphs are influenced by Gestalt groupings (Ali & Peebles, 2012). The arrangement of the various graphical units in a display may therefore affect the inferences drawn on the basis of a given display.

A more active role of information representation is the externalization of information. Humans have limited short-term memory capacity, and offloading information to the environment is a common strategy to facilitate memory. Writing down a phone number or grocery list outsources the memory task to a lookup task once the fact is needed. Instead of memorizing a route or relevant landmarks, one can use satellite navigation to be provided with external clues at decision-making moments.

Scaife and Rogers (1996) provide three core properties of external representations, namely *computational offloading*, *re-representation*, and *graphical constraining*. Computational offloading corresponds to the grocery list example and the intersection inference in the rectangle construction example of Kotovsky and Simon (1990). Re-representation describes the inequality caused by different representations of a structurally similar problem. The isomorphic representations of the Chinese Rings puzzle are one example (Kotovsky & Simon, 1990). Arithmetic with decimal numbers is much easier for someone familiar with it than with the Roman numeral system, even though they describe the same structure (Larkin & Simon, 1987). Isomorphic versions of the game *Tic-tac-toe* differ in difficulty (Zhang, 1997). Graphical constraining arises when certain graphical representations constrain inferences in the described domain. Venn diagrams for syllogisms are an example where the graphic representation bounds the inferences by topological constraints (Stenning & Lemon, 2001; Stenning & Tobin, 1997).

Presenting information in combination with real-world manipulations is another way to reduce cognitive load. Slide rules have long been used to facilitate arithmetic operations through the use of logarithmic scale addition. The mental operation is reduced to manipulating the relative position of two scales, and the result can be read. The slide rule uses the spatial representation of mathematical scales to enable additive calculations for more complex arithmetic operations. Addition and subtraction is a simple operation in the real world, as two lengths can be strung together to calculate the combined length. All three characteristics of externalization of information are used as described above. The physical nature of the slide rule allows the numbers to be read without having to hold everything in memory, thus offloading the calculation. It re-represents the operation $a \times b$ according to logarithmic arithmetic to $\log(ab) = \log(a) + \log(b)$. The display of the logarithmic scales on two individually moving objects allows the addition of the two logarithmic scales and limits the reading of the result to a comparison of a mark on one scale with the other. The slide rule here is a technology that works with the human to aid cognition. The technology that does this ranges from a simple pencil to a supercomputer, neither of which do much without a human interacting with them (Salomon et al., 1991).

2.9. Cognitive Invariants

The extensive experience with spatial information from the environment that humans have has led to the use of spatial information as a tool for problem-solving even in domains that are not strictly spatial, as described above. The influence that this spatial knowledge of the world has on problem-solving is then a relevant and interesting question. Shepard (1994) discusses how such a dependence on invariants from the environment might shape cognition, and why attempting to formulate such invariants might be useful for understanding cognition. The principle is based on similar observations in biology based on how we interact with our world during evolution. For any surface-dwelling organism, for example, a twenty-four-hour day-night cycle is constant. Similarly, objects are constant, which explains effects such as perceived movement in the presence of successive flashes of a stimulus. Shepard argues for the existence of a psychological equivalent to these biological effects and the ability to formalize them conclude them:

> Perhaps psychological science need not limit itself to the description of empirical regularities observed in the behaviors of the particular, more or less accidental collection of humans or other animals currently accessible to our direct study. Possibly we can aspire to a science of mind that, by virtue of the evolutionary internalization of universal regularities in the world, partakes of some of the mathematical elegance and generality of theories of that world. The principles that have been most deeply internalized may reflect quite abstract features of the world, based as much (or possibly more) in geometry, probability, and group theory, as in specific, physical facts about concrete, material objects. (Shepard, 1994, p. 26)

Shepard provides three examples for such perceptual-cognitive universals. Namely, the mental transposition of objects, the color of objects, and the representation of the nature of objects. To achieve invariance in a psychological domain, it is necessary to formalize the effects. Shepard advocates doing this through *representational spaces*. A representational space describes the represented domain when abstracted from material constraints, thus a purely mathematical or geometric representation. The key, then, lies in the definition of a representational space that allows the description and formalization of the observed phenomenon. Shepard has a strong background in mental rotation research and is co-author of one of the most widely used mental rotation tests, the Shepard-Metzler cubes (Shepard & Metzler, 1971). The first example of formalization then comes from this area. The argument is that because objects are generally consistent, two representations of an object in different operations are perceived as a rotation of the objects rather than as two separate objects. Mental rotation of mirrored 2D objects can also be achieved by rotation through 3D space. Humans will do this because they have evolved in a 3D world and are therefore easily able to translate outside the representational plane (Shepard, 1984). Shepard further argues that the mental translation of objects follows the rules of kinetic geometry, just as objects are translated along the simplest path. Kinetic geometry describes the mathematical properties of objects and their possible movement. The movement can, for example, be limited by mutual interpenetration or degrees of freedom (e.g. a hinge), but does not take into account the physical laws of mass, so the overhead of a given translation is not considered (Shepard, 1994). Shepard proposes a formalization of this simplest motion by modeling it with geodesic paths. The perceived color of an object is determined by the light that is reflected from its surface to our eyes. Even though the illumination on the surface itself can vary greatly, the perceived color is unchanging and holistic. Moreover, the perceived color could be generated by mixing three primal components of the color space. There are other effects such as adjacent colors (e.g. B. L. Anderson, 2003; Pinna et al., 2001) and lighting (e.g. Brainard & Hurlbert, 2015; Purves et al., 2001) that could alter color perception based on physically measured color values. Shepard's description is based on perceived color, which may have been affected by these effects. The color space can be represented by the three-dimensional space of *hue*, *lightness/brightness*, and *saturation*. Another color space, also three-dimensional, is *light-dark*, *red-green*, and *blue-yellow*. According to Shepard, this color space is learned through interaction with the world and naturally shifting illumination. Species that live in different environments, such as deep-sea fish or nocturnal animals, may rely more exclusively on the brightness scale than on the hue and saturation scales. Shepard's object grouping example explains it from an evolutionary perspective and links it to affordances in the sense of Gibson (1979/1986). The argument follows the recognition that it is evolutionarily desirable to identify objects of a similar kind in a group rather than a new object each time. Every apple on any given tree does not look identical, yet humans (and many if not all other animals) readily perceive these fruits as members of the same species. This goes further than just fruit from a tree, and categorical knowledge is grouped semantically and hierarchically. Such a hierarchy could be that the apple from the tree is a kind of apple, which is a kind of fruit, which is a kind of food. By categorizing the objects we encounter in the world, we are able to apply learned knowledge even though the situations are never identical. Similarly, the notion of affordances (Gibson, 1979/1986) explains how the perception of the environment is guided by the experience of interacting with certain features. The flat surface of a boulder or a stool or a chair allows you to sit because they are exposed to these structures in the environment. As Shepard points out, the ability to generalize objects into categories is not a failure to distinguish objects, but the ability to attest the same affordance to both. For basic kinds or categories, Shepard proposes a model of a *connected region* rather than a *correlation*. The connected region means that any object in a group can be continuously transformed into another object in that category, while at any point along that transformation, the object remains in the category. For example, an apple can be transformed into any other apple. This approach has the advantage over a correlational approach that a sharp drop in category assignment can also occur for similar objects. Biological mimicry (sometimes called Batesian mimicry) exploits this failure of categorization for protection by making a species appear to belong to a category of dangerous species (Pasteur, 1982). A hoverfly, for example, presents itself as a wasp by similar coloring. With learning two objects can be distinguished through the connected region approach while still correlating heavily. In this way, the poisonous fruits can be distinguished from the edible ones. When a newly encountered object is assigned to a particular category, Shepard formalizes the use of Bayesian inference on the representational space of candidate categories. Criticism of the possibility of such psychological invariants is justified on the grounds that psychology does not lend itself to these rule-like assumptions and that contexts are important. While Shepard was intent on finding universal truths rather than observing statistical regularities (see citation above) Chater and Brown (2008) argue that by modeling cognition in a certain specific situation, we learn more about cognitive processes. The question remains whether certain cognitive effects can be explained by invariant universal properties from the environment or by cognitive mechanisms (Chater & Brown, 2008, p. 60). The authors suggest that these questions should not be answered globally, but separately for different aspects of cognition.

2.10. Summary

The work discussed in this chapter shows that spatial cognition is an important aspect of human problem-solving. The skill with which people use spatial information to solve even abstract problems suggests efficient processes. I have reviewed several studies that focus on the use of space in problem-solving and describe the effects of spatial cognition in general.

Spatial information is a central invariant in the environment because we live in a three-dimensional world, and interaction with the environment is determined by spatial constraints in the environment. Shepard (1994) proposes the advantages of understanding the principle underlying cognition from invariant truths. The work discussed in this chapter allows us to show the effects of spatial information on problem-solving, but the level of abstraction is high at the object level. In the following chapter, I will discuss how I proceed to generate a methodology that not only allows the effects of spatial information to be studied at the level of the information type rather than the object but also allows this information to be manipulated.

3. Theoretical Outline

In the previous chapter, I described the state of research on information use in problem-solving in humans. I have discussed cognitive theories and their view of the role of information in problem-solving. I have further described the special role of spatial information for cognitive beings and especially for problem-solving. Finally, I discussed views on the study of problem-solving behavior and suggested foci for the level of abstraction of analyses.

In the following, I will discuss the theoretical basis of the work for this thesis. I will briefly review why spatial information is relevant in problem-solving and explain how I approach the topic based on the invariant spatial structure. I will then propose a world structure that explains how different kinds of information make up the world as a whole. I will connect this world structure with the idea of different levels of abstraction. Spatial information is retrieved at a low level of abstraction. I argue that the use of spatial information can be studied at such a level of abstraction. I will then suggest some challenges to overcome if such an approach is viable. This gives rise to two main questions that I have addressed in my work.

3.1. Spatial Information in Problem-Solving

Spatial information is central to the function of moving agents. Any form of goal-directed movement requires spatial information. Thus, for humans, spatial information is constantly important, and dealing with it is internalized through constant exposure. Many real-world problems obviously require efficient use of space, such as finding your way home or making and pouring a cup of coffee. For some tasks that require spatial information, this may not be so obvious. Consider and putting a letter into an envelope and mailing the letter. It is not a conscious evaluation of the spatial relationships that one makes to achieve this goal, but something that happens automatically. Extensive practice in the use of our extremities has led to such actions becoming automatic. One theory of cognition based on this premise is evolutionary psychology (Tooby & Cosmides, 2005). In general, evolutionary psychology argues that extensive exposure leads to specific purpose mechanisms for processing certain information. The theory stems from the fact that human organs have also evolved as specialized structures rather than multipurpose organs. Similarly, according to evolutionary psychology, cognitive processes evolved into specialized rather than multipurpose processes. In particular, spatial processing is a core ability for human functioning, and therefore highly adapted processes for processing space exist in the human mind. These are special mechanisms for spatial information processing.

Spatial information is ubiquitous and based on immutable truths in the environment, such as the persistence of objects, so that humans have become very efficient in processing by dealing with spatial information. Some abstract tasks are even solved or communicated with spatial information in the form of metaphors, Like "go ahead with a plan" or "things are looking up". The omnipresence of spatial information and human interaction with it makes it a good candidate for something like Shepard's cognitive invariants (Shepard, 1994). From their approach, I take the idea that people have internalized certain interactions and or properties of spatial information due to exposure to the type of information and use it to solve problems. To investigate this, it is necessary to explore what information humans retrieve from the world when solving problems.

3.2. Information Type Level of Abstraction

My goal is to examine the types of information used during problem-solving. To do this, I classify the types of information so that they are mutually exclusive and collectively exhaustive. These different types of information form the basic building blocks for all information collected. For the purposes of this thesis, I will consider the distinction for the domain of problem-solving. But as discussed in the previous chapter, much of everyday functioning can be described as problem-solving, so I contend that this theory has generality. I will follow the idea of different *levels of abstraction* (LoA) to describe a hierarchical approach to analysis. Rather than looking at what objects a problem solver is manipulating or what goals they are trying to achieve, I focus on a lower level of the types of elemental kinds of information they are retrieving. I would consider the type of information to be very low level for the realm of human perception. I assume that this division into parts of information types is possible and can be studied systematically. Humans do not always seem to be consciously aware of what kind of information they are retrieving, as shown by interference effects such as the Stroop and Simon effects. I have left the term information type undefined so far, in the following part, I will discuss the distinction that I use to make information types mutually exclusive from each other and how they fit into and comprise the structure of the world.

Humans retrieve information from a variety of sources and with different content. In terms of Shannon's theory of information transmission, we could say that there are different channels through which we retrieve information. These might be, for example, the senses. The content of what is transmitted may differ between modalities, but the nature of information is also different. Shannon's theory is not concerned with the content of the messages. The source of the kind of information we retrieve are invariants in the world. This is the low LoA where I want to examine information use during problem-solving. The investigation of the source of an information type takes place at the level of the elements and properties of an object, not at the level of the object from which information is retrieved. That is, there are sources of information sources that meaningfully different, and certain types of information contribute to the holistic perception of a scene. This distinction is not to be confused with the modality of perception. This kind of distinction is more fundamental before any interaction. It is about what knowledge the environment provides, what invariant truths there are in the environment. I will describe these invariants in more detail in the following part.

3.2.1. Structure of the World

I describe the world in terms of elements that make sense to the observer, in my case, for a human observer. I take a hierarchical approach. The *world structure* can be seen as a set of *objects* and their *relations*. Each of these objects and relationships is defined by *structural elements*. And the state of the world is the value of all those structural elements at a given moment.

The *world structure* is the highest level of hierarchy that encompasses the entire environment. Within this environment are *objects*, which in turn can contain objects and *relationships* between these objects. As a *world state*, I define the world structure at a certain point in time. For an illustration of the theory, see Figure 3.1. The diagrammatic representation describes only a small interaction. In the figure, we have two objects that have some relationship to each other, and one of the objects contains a child object. This is to show how the theory is hierarchical. For my thesis, this is the structure for an observer. Some elements of the structure of the world may be different for observers with different abilities. Different sensory abilities or different experiences with certain interactions cause differences in the world structure. Before the development of object permanence, for example, the structure of the world for a child is fundamentally

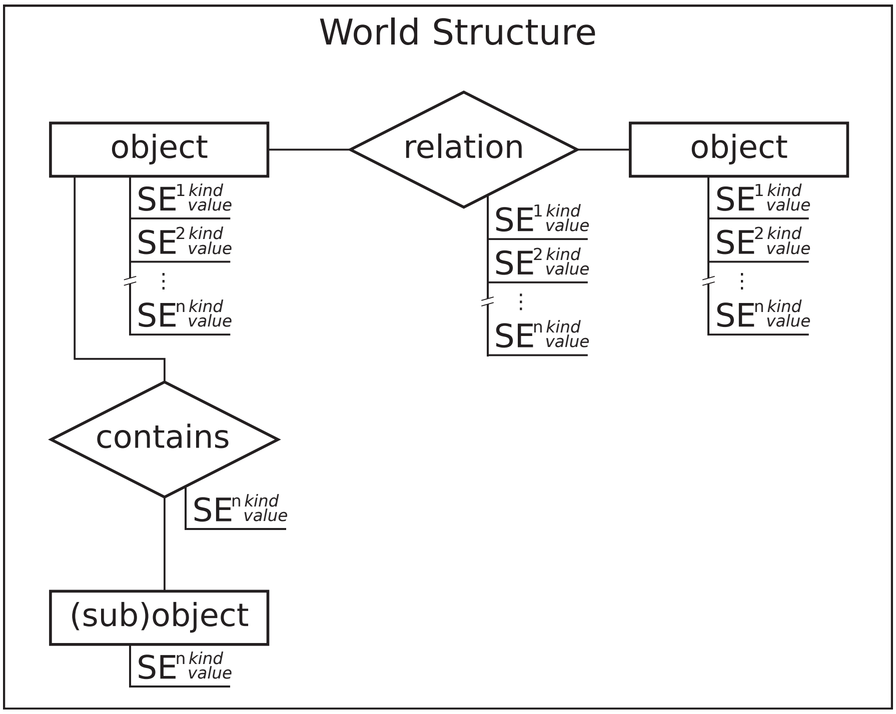

Figure 3.1.: Graphic representation of the world structure model

different from that of post-development.

The structure of the world is determined by the learned invariants it contains, as described by Shepard (1994). The actions available in the world determine how the world can be seen. As with the example of object permanence, we know physical objects cannot float through each other in confusion, so we assume that the object is still under its cover. The structure of the world is therefore dependent on objects and their relationships with each other and ourselves. An object can have any number of properties, such as a specific color, smell, and position. I also assume that objects generally inherit the relationship of their parent object (see Figure 3.1). This can be seen in the cognitive effects described in the previous chapter of directional biases, based on the hierarchical organization of a cognitive map.

For different organisms or agents, the world structure can still be very different. The invariants that are learned differ depending on the interaction with the environment. For a fly, the viscosity of air is relevant and influences its interactions with the world. Because of its low weight, the air offers it valid ways to travel, which may not be available for larger organisms. With a smaller mass in general, the influence of gravity on action decreases, and this is reflected in the structure of the world for this organism. I will focus on man and relate to our world structure when I mention it.

I argue that these relationships and properties can be further described as a set of *structural elements* that are at the lowest level of my hierarchical world structure model. A structural element is a feature of the world structure that is meaningfully different from other elements. I suggest that the meaningful distinction for human cognition is the way that an element is measured. An example of what I consider to be a structural element is color or distance. These structural elements describe the properties of an object and the relationships of the objects to form the entire world structure. The properties of an apple can contain, for example, color as a structural element. Let there be a basket of apples, some red and some green. One apple has a structural element color with a value of red. In this case, the structural element describes a property of the apple. The Gestalt principles predict that a person groups objects by color, so in this case, the relationship between the red apples also contains the structural element color with the value red. Figure 3.1 shows this by the structural elements of different types that describe each object or relationship.

The type of information that the observer retrieves from the environment is based on which structural element is used as the source of the retrieval. When assessing the distance between two objects, the distance information is retrieved by examining the structural element of distance in the relationship between the two objects. Likewise, if you want to assess the height of an object, the structural element of the distance of the boundaries of the object is the source of information. The type of information retrieved from a structural element, I refer to as the *information type*.

Processing these structural elements is unconscious, and information from multiple objects can be retrieved simultaneously. There could be special mechanisms in the perception apparatus to deal with each structural element. Like the color perception in the retina, but also more conceptual elements like topology later in processing, these different streams of information are then combined in the mind to form the holistic representation of one's environment.

The states of the structural elements of each object determine its affordability (Gibson, 1979/1986). The height of a step can afford to step up or to climb up. The flat surface of a stone can afford to sit, while one with a jagged, spiky boundary cannot.

3.2.2. Information type level of abstraction analysis

The *spatial structure* as the part of the world structure, which deals with spatial structural elements, is a prominent source of spatial information. I intend to investigate the use of information in problem-solving by studying the use of information types, in particular spatial information types. For brevity and because the theoretical implications should also apply to nonspatial types, I will use information types without the "spatial". For my thesis, I will focus on spatial information types.

As described above, each system can be examined at different *levels of abstraction*(LoA). For my thesis, the functional lowest level of abstraction is the structural element. An object in the world could be analyzed at an atomic level of abstraction, but as far as the structure is concerned, the structural elements are still the smallest element. For atoms, this could be spatial structural element, but also other structural elements of subatomic particle attractions. To investigate the use of structural elements and various types of information by problem solvers, the LoA should reflect this. I propose an analysis of a *LoA of information type*. For this purpose, it is relevant to have tools that allow systematic monitoring of the use of information types and to manipulate the states of the relevant structural elements in such a way that the resulting effect can be measured. The available operations in a particular situation have a major impact on how computationally demanding a particular problem is.

Consider the problem of Chinese Rings (Figure 2.1b), which was covered in the previous chapter. Multiple versions that are information equivalent may not be computationally equivalent (H. A. Simon, 1978). It is an interesting question whether a difference in the available spatial structural elements and their states could be the mitigating factor for the difficulty in this situation. Imagine trying to find the cheapest deal for an item from an unsorted table compared to a sorted table. In the sorted table, the only operation needed is to compare the top and bottom entries and see which ones are the lowest, this can therefore be solved by a spatial search at the edges of the table. In the unsorted table, it is necessary to traverse the entire table, remembering the lowest price currently found all the time. The latter is obviously more expensive in terms of calculations. The difficulty of a problem may differ with different representations because the structural elements that correspond to the problem structure in one representation may not be represented in another representation. The sorted table allows the use of spatial SE of relationships in the representation to justify the displayed values. Dependency on the spatial structure in the sorted list and the inferred relationships between elements make it possible to reduce the computational cost of finding the cheapest offer. But what types of information do people retrieve from a scene to solve a problem? In order to investigate this, several challenges need to be addressed. In the following, I will specify four challenges for this theory of *information type LoA analysis*, which I will try to answer in the thesis.

3.3. Transfer of Information Types Between Domains

Humans use the types of information types that are familiar to them and that offer the simplest solution for the process in question. Again, the example of the sorted table, where the reader would use the spatial characteristics of the table to get the answer instead of checking each individual item in the table. This is a satisficing approach that is prone to sorting only tables that appear sorted at first inspection but are not, but experience has taught us to evaluate a table and determine if it is sorted quickly. This accepts the possible error that the table was not sorted after all. The term "satisficing" was coined by H. A. Simon (1955) in his theory of bounded rationality, discussed in the previous chapter. In short, it describes the behavior of using an answer without definitely knowing that answer is correct when there is enough evidence that to provide a sufficiently high probability that it is the correct answer. This strategy allows optimizing resources in such a way that after finding a sufficiently supported answer, the pay-off of the effort to continue searching for another answer is too low.

If we assume that human cognition has special purpose mechanisms for spatial information processing, it is likely those will also be used as tools in novel situations or in situations where their use is more efficient than other operations. As with the sorted table example, children also use spatial processing when evaluating the numerical values through number lines (Dehaene, 2001). By the use of spatial processes, i.e., for example, by lining up the numbers from left to right and assigning semantic values of high and low to each extreme, the process of comparing numbers becomes one of shifting attention from one to the other and observing the direction. In this way, the mental load of comparing is reduced to an automatic shift of attention. In this way, relationships in one domain can be used to understand relationships in another domain. For example, on number lines, the positions on the line can be assigned to the magnitude of the value of a number. The relationships of these items can then be used to reason about the relations between numbers. This is a key process by which one learns new concepts, analogy formation (Gentner & Smith, 2013). Analogy building is not the only method of learning, of course, but it is central to the transfer of knowledge and the application of what is learned when faced with a novel challenge. Spatial metaphors like "markets are rising" or "looking forward to something" are often used to describe abstracts concepts of finance and time. Gentner (2001) describes how spatial metaphors form the basis for our communicating about time, which would otherwise be too abstract to be communicated efficiently by other means. Spatial structure and spatial information types are thus important for human thinking and communication about our environment. Thus, for the information type LoA approach, it should be crucial to show that certain information types are used in another domain to support problem-solving. This should be done in a way that allows individual types of information to be isolated in order to assess their use and effectiveness. To prove the use of an information type in a different domain, it is further necessary to be able to manipulate the information source in such a way that the behavioral differences in the solution represent this change.

The first challenge for the information type LoA analysis, then, is a way to study the cross-domain use of information types in a controlled environment.

3.4. Spontaneous Use of Spatial Information

As shown in the previous chapter, people use spatial processes and terminology when reasoning about non-spatial and/or abstract domains. In perceptional processes, the spatial structure of the environment and especially the respective state of the spatial structure influence how a scene is perceived. Object groupings, as described by the Gestalt laws of grouping for an overview, see Wagemans, Elder, et al., 2012, are a clear example of how spatial structure affects perception at a very low level of processing. This interference happens automatically during perception when the scene is interpreted but seems to be later than a purely retinal process, at least at some stage of preprocessing (Palmer et al., 2003).

Consider the task of comparing the height of two fence posts to determine that both were planted at the same height. You could go to the first post and measure the height, remember the value, then measure the second post, and compare the values. Or you could position yourself so that the fence posts are in line, lower your eyes to the top of the fence post, and solve the task by using the top of the post as a sight, taking advantage of perspective and thus spatial structure.

The question remains whether problem solvers, when confronted with a novel situation, actually seek spatial information to help solve that novel problem, even when other means of resolution are more prominently available. Furthermore, this could also lead to detrimental effects on problem-solving behavior, as the spatial structure could be misleading in terms of problem structure.

In terms of LoA information type analysis, the challenge is to have a methodology that allows having a spatial structure that could help or hinder the task at hand, while otherwise being completely independent of the problem. Thus, differences in spatial structure could be manipulated while the problem remained otherwise identical.

3.5. Salience of spatial information

The fact that spatial structures are used in situations where other structural elements are also available could also indicate that spatial sources of information are simply more salient. People can naturally reason with a variety of information types. So could a preference for spatial information be rooted in an inherent difference in salience between structural elements? The grouping of objects according to the Gestalt principles happens automatically during scene interpretation. But does this also apply to other types of information? Just compare the Stroop effect (Stroop, 1935). In this test, the task is to name the color of the ink used to print each word in a list. The words are printed in different colors. The classic effect occurs when the words that are printed name one color but are printed in a different color. Suddenly, the automatic word recognition interferes with the naming of the color because it cannot be ignored.

The salience of spatial structural elements as sources of information for problem-solving is a crucial aspect of the information type LoA analysis approach. The ability to judge the use of structural elements regardless of their importance in a scene must therefore be ensured. If spatial structural elements are simply more salient than non-spatial elements, this could explain humans' use of spatial information in different non-spatial domains. Higher salience simply makes it the first type of information used in representing and solving a problem. A similar case would be that the processing mechanisms for spatial information in humans are so efficient that the use of these mechanisms to solve a task justifies the transformation from another type of information. This case might be indistinguishable from the salience reasoning in some tasks, so the careful design is warranted to have explanatory power in the experimental design. The challenge that arises is that a task must be developed that allows the relative salience of structural elements to be tested when used to solve problems without the likelihood that simple processing will cause a particular structural element to be used.

3.6. Information Type-Based Problem-Solving

The proposed theory of information type LoA analysis is only valid if a reasoning approach based on information types were possible at all. Empirical measures are not sufficient to prove that such an approach could work in a real scenario. Empirical approaches only describe what happens, but not why something happens. This is something that only a synthesizing or modeling approach can provide (Kuipers, 2000). It would therefore be necessary to design a model based on LoA analysis-based reasoning of information type. If it were possible to design a computational model such that it would be able to solve a task by making its decisions based on LoA analysis of information type, it would be much more likely that this would be a cognitively plausible theory of human reasoning behavior. So the challenge is to create a model that can perform one of the designed tasks. In this way, the performance of the model can be compared with the human participants. It should be noted that the goal should be to create a proof of concept for reasoning at an information type LoA rather than to model exactly what people do. The goal of a synthesis approach may be to model human cognition, but at a more fundamental level, it is also a very good way to assess possible problem-solving processes. In this way, it is possible to assess whether a possible approach is at all feasible. A correlation to human behavioral outcomes then provides an indication that the proposed approach might also be how human problem solvers solve the task.

4. Implementation of New Paradigms for Information Type Studies

In this chapter, I describe the steps I took to develop methods that could test the questions raised in the previous chapter. From the theoretical outline, I have defined the following challenges in asserting that reasoning can be studied at the granularity of information types. The experimental paradigms should allow (a) to study reasoning by applying an information type to another domain, (b) to test the spontaneous use of information types and, moreover, also to test whether these information types could impede reasoning, (c) to compare the use of different information types with each other, in addition (d) to carry out such an experiment on a computer model that reasons exclusively on the basis of information types.

4.1. Isolating Information Types

In order to investigate the proposed challenges, paradigms are needed that allow the study of the effect of the use of information types. For experimental designs, this means that one should have the information type as an independent variable. If the information type cannot be used as an independent variable, the behavioral measures must be distinguishable as to which information type the reasoner is basing his decision on. Each of the experiments therefore uses a method of isolating information types to address one of the proposed challenges. To add validity to the proposed study of information use in problem-solving at a granularity level of information types, I developed several novel experimental tasks, each addressing one of the challenges proposed above. I chose to develop different paradigms for each challenge to meet the requirements of that challenge. By designing multiple paradigms, I simultaneously provide a much broader study of the effect of using information types on reasoning. Since there is not much research on this topic, I have therefore developed several sample studies manipulating and examining the use of information types.

4.2. Cross-Domain Reasoning

In the first experimental paradigm, I developed the claim that people can use spatial information to think in other, non-spatial domains. People use spatial metaphors in non-spatial domains such as time, but the goal of this experiment was to determine that participants could also readily use spatial information to reason in a new non-spatial domain. Moreover, the paradigm should also be suitable to test reasoning with other non-spatial information types in a different domain. The spatial information experiment I eventually conducted can therefore be seen as a pilot to confirm the operation of the experimental paradigm. In order to use information from one information type for reasoning in another domain, it is necessary to specify how information from the base domain can be used in another target domain. A common structure or relationships between elements of one domain must be applicable to the other domain. A base domain that shares relationships with a target domain is the formal structure of an analogy. An analogy is the comparison of relationships between elements. Analogies are not literal similes but deal with relations.

As an example, consider this description of malicious computer software. "The program worked like a Trojan horse. It seemed to be a bargain utility, but once installed, it performed an attack on the system." The analogy does not describe the actual appearance of the software, but rather highlights the similarities in relationships between the elements. It highlights the functional role that the software and the mythological, gifted wooden horseplay as an infiltrating element that breaches the defenses of a city or system. So the analogy is that the Trojan Horse is to Troy as the malware is to the computer system. The shorthand notation for analogies, which I will also use in this thesis, uses colons to symbolize the relationships. The above analogy in this notation would be written as follows, Trojan horse : Troy :: Malware : Computer System.

Analogies play a major role in learning and transferring knowledge from one domain to another (e.g. Gentner & Smith, 2013). Analogies are also used as a tool to describe abstract or not so easily perceived concepts through their similarities to physical phenomena. Against this background, analogies offer a great opportunity to be used in a cross-domain reasoning paradigm. People have experience using them, and they are inherently more concerned with relationships than objects, making them suitable for cross-domain use.

In order to reason with information types in another domain, there must be a relationship that allows valid reasoning in the new domain. It is these relationships that are to be conveyed in the analogies. I chose to use *order* as the relationship between elements. Order can be conveyed in all continuous information areas and is therefore ideally suited for use in different types of information. In spatial information, order can be available through different means in different types of information. In topology, for example, an inherent order is given by the concept of a conceptual neighborhood (Freksa, 1991). Conceptual neighbors are those states that can be transformed into each other without an intervening third state. Consider a car parked on a property. If we consider the states of the RCC-8 (Randell et al., 1992), an example of conceptual neighbors would be disconnected (DC) and externally connected (EC). A transformation, for example, but not limited to shifting or scaling, would change the topology between two regions from the first to the second state. On the other hand, disconnected (DC) and partial overlap (PO) would not be considered conceptual neighbors since the transformation of the regions from one state to the other would necessarily result in the state of EC intermediately. Thus, when considering the conceptual neighborhoods, an inherent ordering of the topological states emerges. If the three configurations DC, EC, and PO were represented in their inherent order defined by the conceptual neighborhood, DC could be transformed into EC, which could be transformed into PO. Spatial information in the form of distance information has more obvious inherent order because the scale at which distance is measured is a continuum, so short, medium, long elements could be uniquely ordered according to this criterion. Quite obvious is also the inherent order through directional information in spatial configurations. So an inherent order of three elements could be left, middle, right, or similarly, top, middle, bottom. This order is much more prominent in the environment and is often used to convey information, for example, sorted tables. In the non-spatial domain, for example, order consists of structural elements measured in continuous scales such as pitch, color saturation, and brightness.

A prominent paradigm for reasoning research is the so-called three-term series problem, which consists of two premises followed by a question about the structure posed. For example: "Paul is smarter than Kyle", "Chris is smarter than Paul", "Is Kyle smarter than Chris?" From a logical point of view, these problems do not allow any clear derivations. Rather, their domain is enriched by the solver (Johnson-Laird, 1972). In the example, it is inferred that *smarter* is transitive to arrive at an unambiguous answer. In spatial cognition research, this paradigm is used in directional reasoning tasks. For example: "London is west of Berlin", "London is east of Toronto", "What direction is Toronto as seen from Berlin?" In this case, the directional cues are cardinal directions, and the elements are actual cities, but the same paradigm applies to abstract stimuli or a relative frame of reference. The use of abstract elements in the form of letters additionally eliminates the implicit information of distance that could be formed from prior knowledge. If the reasoner knows where the given cities are located on the globe, the distance between London and Toronto could be shown to be greater than between Berlin and London. Since I want to isolate information types in my experimental paradigms, I have chosen a paradigm with three term-series problems and abstract elements. I have, however, retained the cardinal directions. In my experiment, I used letters as abstract elements.

I used the three-term series problem as the base domain of an analogy and defined a novel abstract target domain that also featured an inherent order but was not based on the information type of direction. For this target domain, I have designed pictograms that contain an inherent order. These were the original designs for this study. Six versions of each pictogram were created, each differing in one characteristic. These versions were so that the pictogram sets contained a natural order. For example, I created a circle with little dots in it. Each version had a different number of dots inside. I used powers of two so that the difference between the circles could be clearly seen. The first circle pictogram contained two dots, the next four until the sixth circle contained 64 dots. Other pictograms were stars that differed in the number of dots, a pie chart that was progressively shaded black, a vertical rectangle that was progressively filled black, a square with vertical stripes that differed in width of stripes, and a square with diagonal stripes that differed in distance between stripes. To form a well-defined analogy from the spatially ordered configuration to these pictograms meant to define the direction in which the pictograms were to be interpreted. This is naturally open to the interpretation of the perceiver. Consider the spatial configuration defined by "A is west of B" and "B is west of C", i.e. A B C. If the basis for the analogy were only on elements of these, the correct answer would depend on the order perceive by the solver in the pictograms. For example, consider the analogy B:C::Pic4:?. Obviously, all other pictograms would be acceptable answers since the direction is underdefined, and distance is not defined at all. Therefore, I used analogies with six terms, each using two relations in the base and target domain, respectively, to establish the frame of reference for the pictograms. So the analogies would be as follows: A:B:C::Pic4:Pic5:?. In this version, Pic4 and Pic5 are set to A and B, respectively, so that the direction is clearly defined, and the only acceptable answer remains Pic6.

As mentioned above, there is no mention of distance in the base section and thus only order should be the basis of the analogy. It follows that for the analogy A:B:C::Pic1:Pic2:?, all other pictograms would be acceptable answers. From the theory of preferred mental models e.g. Jahn et al., 2005 and the fact that in the present experimental design the distance was made explicit by the participant, I expected that participants would consider only one answer acceptable. In the second phase of the experiment, I further made the distance explicit in the base domain and observed whether the participants' response behavior change to integrate the distance information into the representation and change their responses accordingly.

This novel design leverages the human ability to form analogies to enable the testing of cross-domain reasoning at a granularity level of information types. However, the experimental setup is based on making explicit the information to be used in reasoning to arrive at a solution, the following experimental paradigm addresses this limitation.

4.3. Implicit Information Type Representation in Case of Different Problem Structure Correspondence

The first method allows you to test whether information types of one sort can be used in another domain. This allows us to test the ability of humans to reason with a particular type of information – in this case, space – in a different domain. I have described above the special role that spatial information has in human functioning. Humans use this, for example, to aid reasoning about abstract domains. I argue that human spatial cognition is highly efficient and well-trained due to the ubiquity of spatial information. The second paradigm I designed, then, was to test whether people seek spatial information to support problem-solving when the task does not require it. The paradigm was additionally designed to test whether implicitly presented spatial information could also be an impediment to problem-solving.

4.3.1. Requirements for the paradigm

The goal was to design a paradigm that allows testing the effect of implicitly presented spatial information on problem-solving performance. For this purpose, several requirements must be met by the selected method. The method must enable a task that does not rely on the type of information that can be changed. The configuration of the changed information type may be such that the use of that information type is consistent, neutral, or in conflict with the problem structure. This allows you to specify whether this information is used by problem solvers to assist in problem resolution. The difficulty of the task should be such that differences caused by the different configurations can be identified – both facilitation and hindrance.

4.3.2. Game playing as a structured problem-solving task

Games are a form of problem-solving that offers great advantages as a research paradigm. To reiterate: I define problem-solving as working toward a goal using (a sequence of) steps. In the context of games, this very broad definition becomes very tangible. Games have a goal state that is clearly defined, and a concrete rule set that specifies the valid operations that can be performed to achieve that goal. This allows the researcher a great deal of control over the problem space and a systematic study of variations in the game. For the present question, it is important that changes in the information type are implicit. They should not affect the goal, rules, or operations of the game. Instead of changing the rules or goal state, I needed a method to systematically change only the spatial configuration of the game representation.

4.3.3. Tic-Tac-Toe isomorphs

The game I chose for this experimental paradigm is tic-tac-toe, an old classroom favorite to pass time. The traditional problem is as follows. The game is a two-player game. The game board is a three-by-three grid. Players take turns claiming cells in the grid – traditionally, one player marks their cells with an X, while the other uses an O. The goal is to have claimed three cells in a row, column, or diagonal. If the game is played perfectly by both players, it always ends in a draw. The game is fairly simple in the sense that it has simple rules and very limited problem space but offers enough complexity for new players to make it interesting to research. Zhang (1997) used tic-tac-toe as the basis for their experiments on representational effects on problem-solving strategies. The aim of this study was to differentiate problem-solving behavior based on information that is availably represented (external information) and those that must be retrieved from memory or computed (internal information). To test this, the author defined the abstract structure of tic-tac-toe. The *abstract structure* of a problem is one that is independent of the representation of the problem. They then created four structural isomorphs of tic-tac-toe (see Figure 4.1), all of which retained the abstract structure of the game. The abstract structure of tic-tac-toe, as defined by Zhang, consists of four attributes. (1) The game consists of nine elements. (2) There are eight winning triplets, each consisting of three distinct elements. (3) The nine elements are divided into three symmetry categories so that the selection of one element within one symmetry category is indifferent to all others within the same group. (4) Elements in a symmetry category are part of the same number of winning triplets, namely two, three, or four triplets. In the original tic-tac-toe, the attributes are as follows. The nine elements are

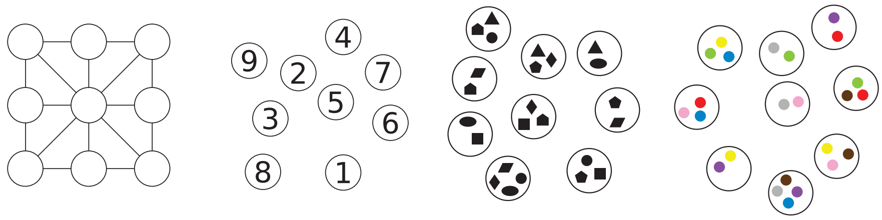

(a) line (b) number (c) shape (d) color

Figure 4.1.: Four tic-tac-toe isomorphs(Zhang, 1997)

the cells of the grid. The winning triplets are the cells in the three rows and columns plus the two diagonals. The symmetry categories consist of the cells at the sides, the corner cells, and the middle cell. Cells in these categories are part of two, three, and four winning triplets, respectively. For example, each corner cell is part of a row, column, and diagonal, so three winning triplets. The isomorphs in Zhang (1997) differed from classical tic-tac-toe in terms of what constitutes a winning triplet. However, the game's play was the same, players took turns claiming elements (circles) until one had claimed a winning triplet, or all elements were chosen without any player claiming a winning triplet, resulting in a tie. So the instructions differed in what a winning triplet was. With the line variant (Fig. 4.1a), the goal is to claim three elements connected by a straight line. With the number variant (Fig. 4.1b), the goal is to claim three circles whose numbers add up to 15. With the shape and color variants (Fig. 4.1c & 4.1d), the goal is to claim three circles that have the same shape or contain the same colored dot. To

(a) line (b) number (c) shape (d) color

Figure 4.2.: Tic-tac-toe isomorph comparison of problem structure(Zhang, 1997)

illustrate the structural isomorphism of these variants, see Figure 4.2. In this figure, the spatial configuration of each variant has been modified to represent that of the classic tic-tac-toe. It was obvious that each variant had nine elements, it is now clear that each can be reduced to classic tic-tac-toe. The winning triplets are the rows, columns, and diagonals in the transformed representations. The symmetry categories are also the elements on the sides, in the corners, and in the middle. The number of winning triplets to which each element belongs is also the same as the classic tic-tac-toe cell in the same place. This property is reflected differently in each variant. For the line variant, it is the number of lines that intersect a circle. For the shape and color versions, it is the number of shapes/dots in each circle. In the number variant, it is the even numbers, the odd numbers, and the five. In this variant, the representation has no explicit graphical clue as to which elements belong into the symmetry categories, as with the number of dots in the color variant. The symmetry groups must therefore be calculated/known for this variant.

The experimental task in the study of Zhang (1997) was a series of games of one of these variants against a computer opponent. The task was the same for each variant. The computer player would always choose first and play perfectly. This meant the optimal outcome for the participants was to tie the computer. Participants were informed of this and that there is a strategy in place to ensure that they consistently tie the computer. Their task was to find this strategy. The experiment ran until the participant had tied the computer ten times in a row, assuming that the player had found the successful strategy. If the participant had not tied the computer ten times in a row by the end of the 50th game, the experiment was also stopped.

Zhang (1997) found that participants who played the line variant performed significantly better than participants who played the other variants and that participants performed averagely in the color condition, while participants performed poorly in shape and number conditions. Zhang has studied problem-solving behavior quantitatively and qualitatively. Since they were dealing with strategy changes, a large part of their analysis is devoted to per-move decisions, which was not a focus of my work. Instead, I focused on the performance measure they used as a quantitative measure, namely the percentage of players who successfully found a strategy to consistently tie the computer (*success ratio*) and the number of games it took a participant to find that strategy(*games before strategy*). The success rate was determined per condition (variant) as the percentage of participants that achieved ten ties in a row. Pre-strategy games were defined as the number of games up to and including the first game of the ten-game tie streak. The success rates per condition were, 100% in the line variant, 45% in the number and shape variants, and 80% in the color variant. The average games before the strategy were 6.5 in the line, 40.5 in the number, 39.5 in the shape, and 27.8 in the color variants. Participants who did not find a strategy within 50 games were counted as having 50 games before strategy[1]. Zhang (1997) found that the difficulty of the problem of finding the

[1]This leads to a ceiling artifact in the results and thus the performance of the shape, number, and color

correct strategy differs significantly between different representations of the game. The simplest was the line variant, which was essentially the original tic-tac-toe. This was also the only one where the spatial structure of the representation matched the problem structure.

Given the strategy used by the computer, which was only one of several perfect strategies, the player's first two decisions are crucial for achieving a draw. For this illustration of the strategy, consider the classical tic-tac-toe, which works analogously in the isomorphs. The computer starts with a corner element. The player must select the middle element. The computer then takes the opposing corner as its first choice. The player must now select a side element. If the player takes a corner, the computer takes the remaining corner and has created a fork – a situation where there are two possibilities for a resulting triplet, only one of which the player can block one with the next move. The analysis of the participants' problem-solving behavior showed that in the shape and color variants, participants used the heuristic more-is-better in selecting the elements – they chose those elements that are part of the highest number of winning triplets. This succeeds in the first step with the color and shape variant, which claims the item with four objects inside. In the number version, a bigger-is-better heuristic already fails on the first move because you should choose the "5". The computer's strategy was specifically designed to expose the fallacy of the more-is-better heuristic on the player's second turn when the correct choice is a two-object item rather than one of the three-object items. The number version is a little different, here the higher-is-better heuristic would actually lead to a correct choice("9") since the computer would have chosen two even numbers, but most participants made this mistake on the first step if they followed the higher-is-better heuristic. The more-is-better heuristic could also be applied in the line variant, which is essentially classic tic-tac-toe, but it appears that participants were able to correct their behavior more easily in this condition.

In a follow-up experiment in the same paper, the author investigates the effect of spatial symmetry of the representation on the difficulty of the problem (Zhang, 1997, experiment 3). For this purpose, the line variants were modified in such a way that the symmetry groups of tic-tac-toe were also spatially symmetrical or not. The author found that the performance increased significantly when the spatial symmetry matched

variants are likely overestimated. The experimental design did not allow for a participant to record a successful strategy found after game 41 as the experiment would always stop at 50 games and thus ten draws could not be reached. This means that some of those who did not find a strategy may have actually found one within the last ten games, so rating them as 50 would underestimate their score. Similarly, those who actually never found a strategy and scored 50, which overestimates their performance. In my experiments, this was addressed as described below.

the symmetry groups of the problem. This strongly suggests that the spatial structure of a problem representation can support problem-solving. What has not been explored is whether this spontaneous use of spatial structure as a tool might also be detrimental to problem-solving.

There are effects of interference of irrelevant spatial information in stimulus-response paradigms. For example, the Simon effect, in which responses are faster when they occur on the same egocentric side as the stimulus presentation (for a overview, see J. R. Simon, 1990). This is also the case with interference in the Eriksen flanker task, where irrelevant distractors flank the stimuli and provide conflicting information about the stimuli (Eriksen & Eriksen, 1974). It has been shown that spatial interference also occurs with interfering spatial cues in the form of arrows (e.g. Kopp et al., 1996). In these paradigms of rapid response, the presence of conflicting information requires quick interpretation of the available information. The question remains whether such interference can also occur in more complex, problem-solving tasks.

To test this, I adapted the color variant of Zhang (1997). I chose this option because the performance in this state was average. This allowed us to test facilitation and obstruction of the spatial structure of the problem representation in one experimental paradigm. The average performance additionally reduces the chance of confounding the data with ceiling or floor effects.

I have created three variations of the color variant (Fig. 4.1d) for my experimental paradigm. The rules of the game have remained the same, and I have also kept the task of playing against a computer opponent. This allows for a task that is not inherently spatial – the instructions stated the goal as an assertion of three circles that share a dot of the same color. The spatial information was not necessary to solve the task. The change between the three variants I created was spatial, thus allowing us to examine whether problem solvers used spatial structure to support their problem-solving process. The spatial configurations of the three variants were either consistent, neutral, or misleading with respect to the problem structure. As shown above, the color variant of tic-tac-toe can be reduced to the classical tic-tac-toe by rearranging the spatial configuration of the elements. This type of representation has a spatial structure that matches the problem structure. The number of elements is more easily derived by perceiving the three-by-three grid as one group according to Gestalt grouping principles (for an overview, see Wagemans, Elder, et al., 2012). The eight winning triplets correspond to the rows and columns, as well as the diagonals, which are also a natural sub-grouping of the game board according to Gestalt grouping principles. The three conceptual symmetry groups are also spatially symmetric in this representation. And the interchangeability of these groups is made clear by the rotational symmetry of the game board. That is, the problem can also be solved by completely ignoring the color information and solving the inherently spatial classic tic-tac-toe. The line variant in the study by Zhang (1997) most closely represented this configuration, and I expected the performance in this category to be the best of the three. The neutral spatial configuration was the same as that of the color variant in the original study. The configuration and distribution of the elements were pseudo-random so that the spatial relationships between the elements had no correlation to the problem structure. In this condition, the participant would not find any help in the spatial configuration to solve the problem. The third variant I did was the misleading one. Here the arrangement of the elements was similar to the consistent variant in a three-by-three grid, but the spatial structure did not match the problem structure. The elements were distributed within the grid in such a way that reliance on the spatial structure as a problem-solving aid was rather detrimental to problem-solving performance. In contrast to condition two, the spatial configuration implies a natural grouping of elements. I expected the performance in the third condition to be the worst of the three. An illustration of the three conditions can be seen in Figure 4.3.

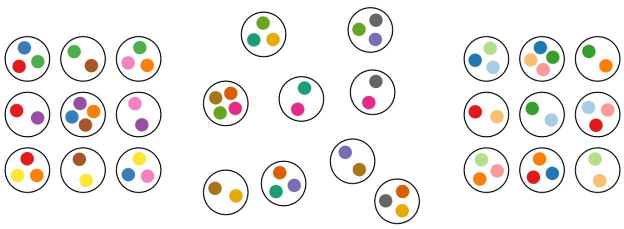

(a) Consistent (b) Random (c) Conflicting

Figure 4.3.: Example configurations of each condition of the experiment.

4.4. Competitive information type use

The third challenge was to differentiate the use of information types, i.e. to create a methodology to test readiness to use different types of information against each other. This was done to assert that spatial information is not simply more salient than other types of information and used in problem-solving. This means a task is needed that allows different information types to compete with each other, and that generates behavioral data indicating which information type is used and which information types are generally more common.

The task should provide information of different information types simultaneously to reduce sequencing effects. The task should involve decision-making. The behavior of the solver should reflect what type of information was used as the basis for the decision that led to the given behavior. Of course, the task should not differ between information types. To test whether spatial information is more salient than other types of information, it is necessary that both spatial information and non-spatial information can be represented. This led to the choice of using graphic stimuli, as this medium makes it relatively easy to represent spatial and non-spatial elements simultaneously. This also facilitated the presentation of the stimuli in an experimental setting. Nonverbal reasoning tasks in tests of mental capacity such as IQ tests have long use graphical representations to address language limitations. Raven's progressive matrices (for an overview, see Raven, 2000), for example, use graphical representations with stimuli that vary in different properties across multiple information types.
Another area where information of various kinds is relevant is matchmaking games such as rummy games. Here, the sets must match in one characteristic – the card value – while differing on another – the suit. The other valid grouping is runs that match in suit but have different consecutive card values. A more recent game that has taken this matchmaking further is the card game SET[2]. The game consists of cards on which there are symbols with four characteristics – number, shape, color, filling. Each characteristic has three states. There can be one, two, or three shapes on the card. The shapes can be an oval, a diamond, or a squiggle. The shapes can be printed in red, green, or purple. The shapes can either have just an outline, be completely filled, or have a striped filling. Each card in the deck has a combination of the four characteristics on it, so there are $3^4 = 81$ cards in the deck. The object of the game is to find sets from a pool of twelve cards. All players watch the same board. Once a player finds a set, he declares it and collects the cards that correspond to the set. The board is then replenished, and the game continues. The winner is the player with the most sets if the board cannot be replenished and no sets remain. A set consists of three cards where each the state of each characteristic state must be either the same or different – i.e. three different states – on all cards. A sample game board is shown in Figure 4.4. In this example, the one possible set would be the green diamond, the green ovals, and the green squiggles. These differ in all characteristics, so that there is only one state in the set, except in color, in which they are identical.

The SET game provides a good starting point for representing different types of in-

[2]SET is a trademark of Set Enterprises, Inc. The SET cards shown in this thesis are merely illustrations by the author and are not original cards.

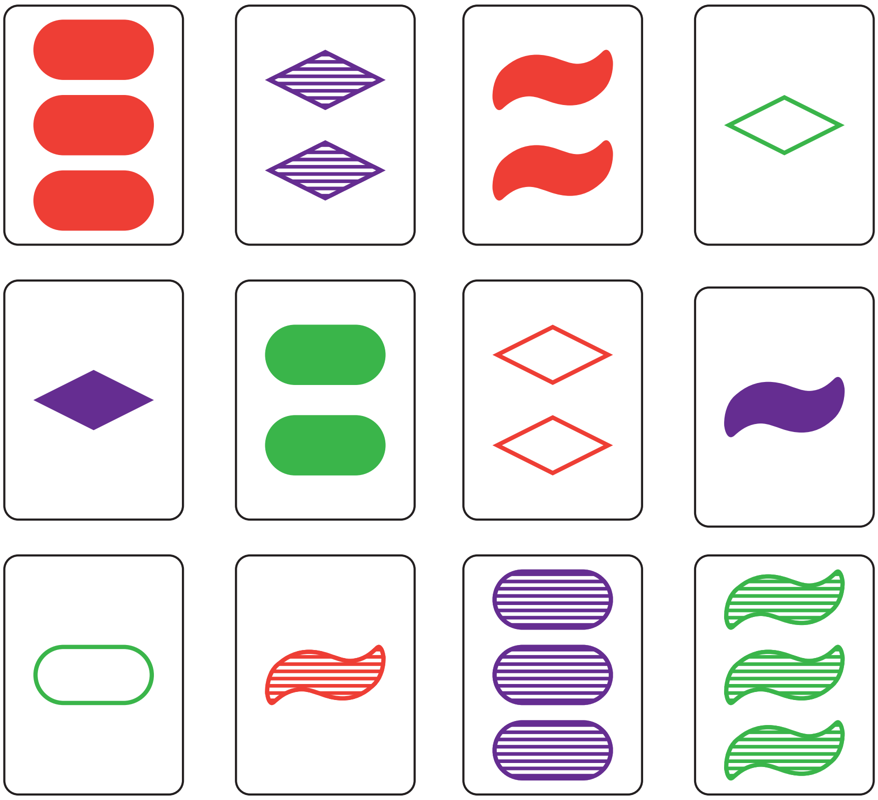

Figure 4.4.: Illustration of a possible board state in the game SET

formation and integrating them into a decision-making process. However, the different characteristics are not clearly spatial or non-spatial. Also, the game rules are prone to search pattern artifacts, and I would have expected a lot of noise. Therefore, I looked for a slightly adapted task, using SET-like stimuli to take advantage of the different types of information.

There is a psychological testing paradigm of sorting based on one of several criteria in the Wisconsin Card Sorting Test (WCST; E. A. Berg, 1948). This test is used to test the ability to adapt to novel situations in problem-solving. It is commonly used in the diagnoses of damage or impairment of the prefrontal lobe (e.g. Demakis, 2003; Nyhus & Barceló, 2009). Patients with lesions of the prefrontal lobe are generally unable to adapt adequately to new situations in problem-solving. The task achieves these novel situations by having the participant switch sorting criteria unnoticed. The test consists of a set of cards similar to the SET cards. As in SET, each card has shapes printed on it that differ in certain characteristics. In contrast to SET, there are only three characteristics, but with four expressions each. The shapes available are stars, crosses, triangles, and circles. Each card is printed in one of four colors, red, green, blue, or yellow. One to four shapes are printed on each card. Of these 64 cards (4^3), four were considered as target cards (also called stimulus cards, I chose to call them target cards for clarity) and replicated. These target cards are chosen so that the state of each characteristic is represented on exactly one card. The target cards are placed in a row in front of the participant. The task is to take one card at a time from the deck and sort it by any characteristic. Without the participant knowing it, a characteristic is set as the sorting criterion. After each sorted card, the experimenter answers whether the sorting was correct according to the established criterion. Note that the participant may receive a false positive if the participant's hypothesized criterion is randomly sorted to the same target card as the actual criterion. When the participant has correctly sorted a certain number of consecutive cards, the experimenter changes the sorting criterion without telling the participant. The participant now receives negative feedback on what was previously current. The sorting strategy must be adapted to any arbitrary change in the rules of the game. This ability to switch strategies appears to be impaired in patients with frontal lobe lesions. Failure to adapt may therefore be an indicator of neurophysiological damage. That is not what made this experimental paradigm interesting to me. The setup offers great potential for testing the importance of different types of information against each other. The cards can be sorted by any of the characteristics, and there are no tasks other than sorting that might influence the sorting decisions. The different types of information are presented implicitly and not given a cue, so their recognition is based purely on salience. The changes between sorting criteria offer an intra-subject repetition of the strategy finding, i.e. to find the right criteria, and therefore repeated measurements of the relative importance of different types of information. Also, during the decision problem, the entire state space of the sorting criteria is visible in the target cards, minimizing the effect of the card presentation order.

As with the SET cards, there are no clear spatial versus non-spatial information types for the cards used in WCST experiments. I therefore developed novel stimuli based on the ideas of SET and the WCST to include both spatial and non-spatial information types. I chose to adopt two non-spatial characteristics of the cards mentioned above, namely the color and the fill. These can be easily adjusted in a graphical experimental environment. The shape of the forms represented was not decisively spatial or non-spatial. One could argue that shape clearly differs only in spatial arrangement and the number of vertices that make up the shape, but since humans perceive objects as a whole and not their parts(Wagemans, Elder, et al., 2012), one can also argue that the shape of symbols is not spatial. Similarly, size can be spatial in the sense that the distance between edges is increased, as is the surface area, but it can also simply be an attribute given to an object that is *large*. I have chosen to use two spatial categories, location and topology. Location is inherently spatial and is not easily confused with other types of information. I chose to represent the experimental stimuli on cards as in the aforementioned paradigms, but on virtual cards, to allow a computer-guided experiment. The different location states were achieved by placing the symbol either at the top, middle, or bottom of the card. The final characteristic I chose was topology. To represent the topology, I used two shapes, a circle and a rectangle. The three states were that the rectangle contains the circle, that the rectangle and circle overlap, and that the rectangle and circle are not connected. From these characteristics, I created 81 cards (3^4), three of which I used as target cards. The target cards, showing all possible states for each characteristic, can be found in Figure 4.5. As a task, I followed fairly

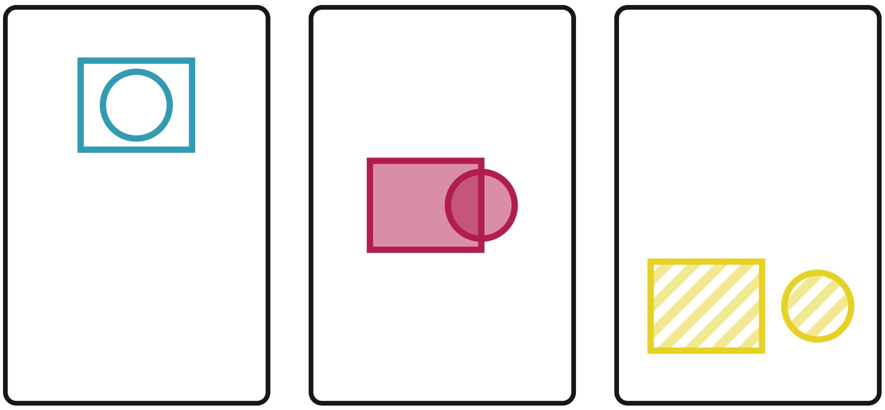

Figure 4.5.: The stimulus cards used in Experiment 3

closely the standard WCST procedure described above. Thus, participants sorted cards and received feedback if the sorting was correct. After the participant learned one criterion, it was switched to another criterion without the participant's knowledge. The sorting criteria changes were designed in such a way that it was possible to compare how easily each characteristic was identified as the initial criterion, but also how the change from one type of information to another affected the relative importance of the characteristics. So if switching from one spatial information type to another spatial one is easier than switching to a non-spatial information type as the sorting criterion? The experiment was conducted online on a computer. This digital version allowed further testing of a proof-of-concept computer model whose conclusions were based solely on the relative importance of information types to accomplish the same task and compare its performance with the empirical data.

4.5. Modeling an Information Type Strategy for a Categorization Task

As a proof of concept for the applicability of information type granularity arguments, I designed a cognitive computer model that performs the card sorting experiment task. The model was able to perform the card sorting experiment in the same way as a participant in the original study. All sorting decisions were based on feedback and the relative saliance of each of the four types of information. I have designed several versions of the model, from very basic to more advanced with more parameters than just the saliency information type. The models were designed with the philosophy of building a very simple mechanism and observing how it would perform in the given task. And to add complexity, only when it represents a significant gain to the performance of the model or correlation to the empirical data. Each model should therefore be able to perceive a card, choose a sorting target, and receive feedback on that decision. The models thus generated were not adaptive in the sense that a model with more experimental runs bar by noise would perform better. In the spirit of keeping the models simple, they were given the ability to perceive card states (information type states), know how to select a target, remember their current sorting hypothesis, and show a response to feedback. Each model run started with set saliency values for the information types, i.e. color, fill, topology, and location. The model would then choose a hypothesis for sorting based on the relative saliance of the information types. If the feedback was positive, the same hypothesis was retained for the next trial. In the case of negative feedback, a hypothesis was again randomly selected according to the relative probabilities given by salience. I then had the model perform the experiment and compared the behavioral results with those of the participants in the card sorting experiment. In this way, a good fit between modeled data and empirical data allows for grounded theories about how the salience of information types might explain the empirical results. I have developed more complex models than the one described above to improve the correlation to the empirical results, a more detailed description of the models used follows in the next chapter.

A cognitive model whose reasoning is based solely on the information type that provides data consistent with empirical data would provide strong support for the theory of information type granularity reasoning. The model can be seen as a proof of concept for a novel way of looking at problem representation and perception. Synthesizing the behavior, rather than just using statistical analysis, gives us an indication as to why the behavior occurs, not just that it occurs. This has been noted as a major advantage of such design approaches over empirical approaches (Kuipers, 1983). Since the model is designed to test the relative salience of information types, I named it Information Type Salience Model (ITSM).

5. Experiments

In this chapter I describe the three experiments I conducted using the paradigms designed in the previous chapter. [1]

5.1. Experiment 1: Inter-domain analogy solving

The first experiment was designed to test the transfer of information types from one domain to another. For this, I used an analogy-building paradigm.

5.1.1. Method

Participants

Twenty participants were recruited through the University of Bremen online bulletin board and received monetary compensation or course credit. The data of five participants were corrupted during recording and could not be analyzed. The data of another participant had to be excluded because he did not follow the instructions of the experiment. The data of 14 participants could thus be analyzed: 7 female; $M_{age} = 26.3$ years; $SD = 8.3$; age range 20–52.

Task

The participants were asked to solve a three-term-series problem. They were successively given premises such as "*A* is west of *B*", "*C* is east of *B*". After reading the premises, the participant was presented with the three letters and had to arrange them to represent the structure described; in the above example, *A B C*. The participant had to remember this configuration, which then had to be used as the basis for an analogy to a continuous range of pictograms (see Figure 5.1). An analogy could be, *A* is to *B* is to *C* as *Pic3* is to *Pic4* is to *?*. Acceptable answers here would be *Pic5* and *Pic6*. The participants had to arrange the given four remaining pictograms in order of acceptance. They

[1] Stimuli and supplementary materials of the experiments designed for this thesis are available at: https://github.com/wienemann/experiments

were also provided with a graphic cut-off line; acceptable answers had to be placed above this line and unacceptable ones below. This response method was adapted from Rumelhart and Abrahamson, 1973

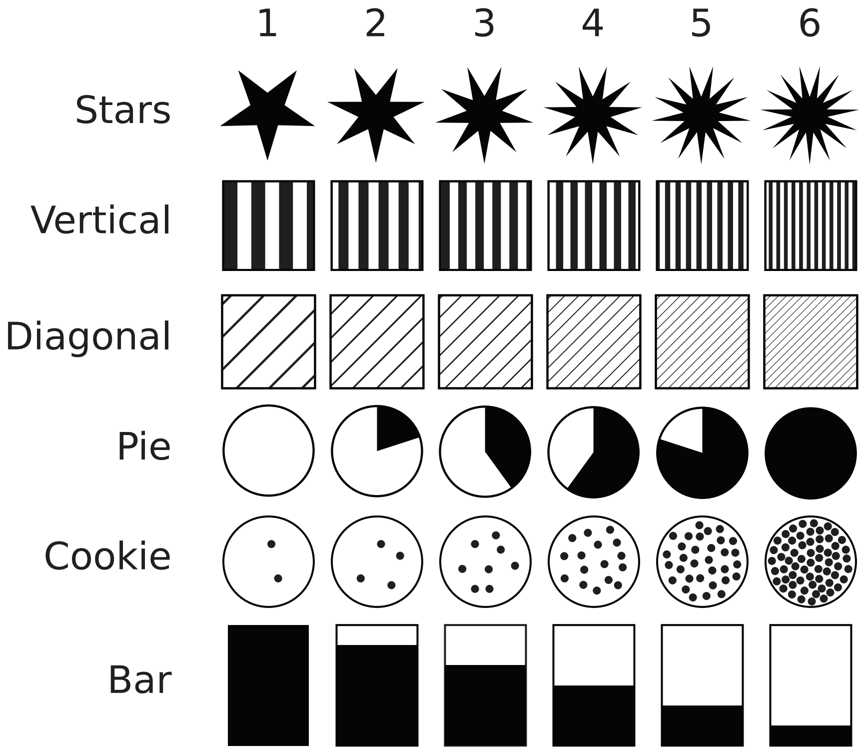

Figure 5.1.: Analogy target pictograms

Materials

Setup. This experiment was performed on a computer and was run in a browser (Firefox v.28.0.1). Interaction with the experiment was mainly with the computer mouse and for the final questionnaire, also with the keyboard. The experiment was conducted in a room with darkened windows to reduce distractions. The data was recorded locally in a text file.

Stimuli. The experiment was presented to the participant in full-screen browser mode (screen size 19inch, resolution 1280 $\times$ 1024 pixels) so that no user interface elements were visible except those used to interact with the experiment. The interactive display area was a 770 $\times$ 770px, with a white area against a gray background. All Instructions and the tasks were presented within this area. Below the area was a button to advance the experiment. The instructions were displayed sequentially, with the participant clicking the button to continue once he or she understood the instructions. The experiment was conducted in German.

The participant had to rearrange letters or pictograms to solve the tasks by drag and drop. This was achieved using the KineticJS (v.4.0.4) library, which extends the HTML5 *canvas* element.

Three-term series problem. The interactive area displayed the premises for the three-term series problem. After both premises were shown and understood, the three letters of the series appeared next to each other in random order. These letters were in a gray "tray" at the upper edge of the display set in "Andale Mono" displayed in 32pt. The letters could be moved in any order to represent the configuration described in the premises. Participants were instructed to place all the letters from the tray onto the interactive area. There were six different problems, and each had its own set of letters as shown in Table 5.1. Two of these six tasks called for a target between the two reference letters and four called for targets at one end of the series. Three of the tasks required a horizontal arrangement with the terms *east (östlich)* and *west (westlich)*, and the other three required a vertical arrangement with *north* and *south*. Items 7 through 12 are replications of the first six items with the addition that the directional cue (e.g. *east*) was accompanied by a distance modifier, either *far away (weit weg)* or *near by (nah dran)*.

Analogy. After the participant had completed the three-term series task, the analogy task could be started by clicking the *next* button. The analogy missing the last term was displayed at the top of the interactive area. The remaining four pictograms of the series not used in the analogy were presented in random order vertically on the left side on a *tray* similar to the three-term series. A line has also been drawn in the middle of the interactive area to indicate the cut-off line for acceptable responses. The pictograms could be moved individually by drag and drop. The pictograms used are shown in Figure 5.1. I designed these pictograms to have a natural order but to be in an abstract domain. The arrangement of the pictograms was validated in a pretest.[2]

[2] $N = 9$ (6 male), $M_{age} = 27.9$ years, $SD = 4.26$, age range 21–33. All participants found an intended order in the *pie*, *bar*, *cookie*, and *diagonal* pictograms. Eight out of nine found an intended order in the *vertical* pictograms, and seven out of nine found an intended ordering in the *star* pictograms.

Table 5.1.: Analogy experiment items

Distance Cue	Item	Configuration	Premise 1[1]	Premise 2[1]	Pictogram	Analogy	Target sequence
Without distance	1	KMX	K w M	X e M	pie	M:X:K::4:3:?	5,6 | 2,1
	2	YBP	B e Y	P e B	bar	Y:P:B::2:6:?	4,5,3 | 1
	3	XRB	X w R	B e R	cookie	X:R:B::3:4:?	5,6 | 2,1
	4	ENF	E n N	N n F	vertical	E:N:F::4:3:?	2,1 | 5,6
	5	STX	S n T	X s T	star	S:X:T::5:1:?	3,4,2 | 6
	6	EGW	G s E	W s G	diagonal	W:G:E::3:4:?	5,6 | 2,1
With distance	7	R MW	R w M f	W e M c	pie	M:W:R::4:3:?	6 | 5,2,1
	8	XZ T	Z e X c	T e Z f	bar	X:T:Z::2:6:?	3 | 4,5,1
	9	CY A	C w Y c	A e Y f	cookie	C:Y:A::3:4:?	6 | 5,2,1
	10	OW J	O n W c	W n J f	vertical	O:W:J::4:3:?	1 | 2,5,6
	11	PK T	P n K c	T s K f	star	P:T:K::5:1:?	4 | 3,2,6
	12	O UE	U s O f	E s U c	diagonal	E:U:O::3:4:?	6 | 5,2,1

Note. Items 2,5,8, and 11 are items that prompt for a response between the reference letters. The remaining items prompt for a response at one end of the continuum. The numbers in the analogies refer to the respective pictogram. The target order is the expected target order and the position of the cut-off line.
[1]n,e,s,w: north, east, south, west respectively. f: far away, c: close by.

Questionnaire. The final questionnaire was also conducted on the same screen as the experiment. It consisted of questions about gender, age, nationality, the experienced difficulty of the experiment on a five-point scale, how the participant had heard of the experiment, and a field for general comments about the experiment.

Procedure

The participants were welcomed and took their seats in front of the experiment computer. After signing the consent form, the participant processed the entire experiment on the computer. The participant first read the instructions on the screen and then solved one training exercise each for the three-term series problems and the analogy. The experiment would not start until the participant had successfully solved both training examples. At this point, the participant was asked if everything was clear, and if so, the experimenter left the room. The participant then solved six trials without an explicit distance cue (no distance condition), followed by six trials with an explicit distance cue (distance condition). Each trial consisted of six steps. At each step, the participant had to click a button within the playing field to move to the next step. The control over the running of the trial was therefore in the hands of the participant. The steps were: (1) The first premise was shown; (2) the second premise was shown; (3) the three letters were presented, and the participant ordered them according to the premises; (4) a prompt to solve the analogy was displayed; (5) the analogy and possible target pictograms were presented, the participant had to rearrange them; (6) a prompt was given to start the next trial. After the completion of the twelve trials, the demographic questionnaire was submitted. After that, the experiment was over, and the participants were thanked and received their compensation.

5.1.2. Experimental design

The experiment was a single-factor within-subject design. The independent measure is the presence of explicit distance information versus no distance information. The dependent measure was the success of finding an acceptable analogy target and the number of targets deemed acceptable.

5.1.3. Results

The data collected for this experiment were analyzed with respect to two main questions. First, in terms of the difficulty of the task and the general ability of the participants to solve the analogies correctly. Second, the answer behavior is analyzed in general and with regard to the preferred answers.

Difficulty of the task

Participants rated their perceived difficulty after the test on a five-point Likert scale from 1 *very easy* (*sehr leicht*) to 5 *very difficult* (*sehr schwer*), $M = 2.93 (SD = 0.83)$. Thus, most participants found the task to be of average difficulty, only one participant rating the task as *very easy*, all others rated in the range of 2–4.

Table 5.2.: Analogy experiment performance

Condition	middle trials	extreme trials	total
no distance	.75	.73	.74
distance	.82	.61	.68
total	.79	.67	

Note. Proportions of successful trials in the Analogy Experiment per condition and type of trial $N = 84$.

As can be seen in Table 5.2, participants successfully completed 71% of the trials. Success was defined as the correct representation of the three term series configuration in Part One of the trial and the formation of a correct analogy in Part Two. Analogies were rated as correct if the participant rated as acceptable only those answers that did not break the structure of the three-term series problem. Depending on item type, this can be either three or two pictograms. Per condition, percentages were similar, no distance 74% and with distance 67%. The difference is probably explained by the additional possible responses in the no-distance condition. Similarly, the differences in performance between intermediate items – the target responses lie between the two pictograms in the analogy – and extreme items – target lies at one end of the spectrum – were close to 79% and 67% success rate, respectively.

Response behavior

First, I looked at how many of the participants actually gave more than one acceptable answer. Again, these are just correct trials neglecting the distance information in both conditions. The results are shown in Table 5.3. Further analysis of the response behavior

Table 5.3.: Analogy experiment number of options given

Options given	No distance	Distance	Total	Percentage
One	47	46	93	78
Two	13	10	23	19
Three	2	1	3	3

Note. Distribution of the number of acceptable options reported by participants per condition.

is divided into moderate and extreme items. The preferred responses per condition and item type are shown in Table 5.4. Responses are categorized as equal distance or edge. The “Equal distance” option assumes that the distance between all three elements is equal. In the moderate items, this describes the middle of the three acceptable options. In the extreme cases, the “Equal distance” is the one that is closer to those given in the analogy. In the moderate cases, it is the edge options that are on the side of the middle option, and in the extreme cases. it is the extreme option at the end of the spectrum. The distribution of preferred responses was analyzed using a Chi-Squared test to determine whether the manipulation of the distance cues had an effect on the response behavior. As expected based on the distribution for both item types, the test showed a significant deviation from chance: For the moderate items $\chi^2(1, N = 44) = 7.29, p = .007$ and the extreme items $\chi^2(1, N = 75) = 27.62, p < .001$.

Table 5.4.: Analogy experiment response distribution

Item type	Equal distance	Edge
Moderate items		
No distance	14	7
Distance	6	17
Extreme items		
No distance	31	10
Distance	5	29

Note. Response distribution per condition and type of item.

5.1.4. Discussion

The results show that the participants were able to perform the cross-domain analogies well and did not find the task too difficult. The success rates confirm this assessment. The translation from spatial structure to a transitive ordered abstract domain did not present a major hurdle for the participants. Furthermore, the manipulation of one type of information was successful. Participants changed their response behavior with the addition of the distance cue. The preference for a single response in the majority of trials is consistent with findings of preferred mental models in reasoning (Knauff et al., 1995; Rauh et al., 2005).

5.2. Experiment 2: Spontaneous Use of Spatial Structure in a Problem

The second experiment examined the question of whether the use of spatial structure found in Experiment 1 was also present in a situation where the spatial information was not necessary for the solution and the use of spatial information was not explicitly mentioned. Furthermore, whether a misleading spatial structure could also hinder problem-solving performance.

5.2.1. Method

Participants

The participants were students of the University of Bremen, who participated for credit and monetary compensation. Each participant received an hourly compensation (€8.50/hour) or course credit if desired/eligible. Each participant also received a bonus based on their performance (€0.05/tied game) to encourage players to find a strategy.

Three participants were excluded from further analysis because they did not follow the instructions, and one participant was excluded because she could not complete the task independently. Thus, all analyses were conducted using data from the remaining 60 participants; 41 female, $M_{age} = 23.4$ years, $SD = 3.5$, age-range 19–44. All participants passed a color vision test[3] and had a normal or corrected-to-normal vision.

Task

The task was the same as in the study by Zhang (1997) and differed fundamentally only in the materials used. Participants were asked to play a game of the tic-tac-toe Isomorph against a computer opponent. The computer would always play first and played perfectly, i.e. only a draw was possible for the participants. The participants were told that they could not win, and the goal was stated to be a tie against the computer. They were also informed that there was a strategy that would allow them to always play a tie against the computer. Each participant was told they would play 50 games against the computer and would receive a bonus of €0.05 for each game tied, a maximum of €2.50. This was done as motivation to find a winning (draw) strategy. In reality, participants could play more than fifty games. In the study by Zhang (1997), the experiment was terminated after 50 games, i.e. from game 42 onward, no winning streak could be achieved as the participant was no longer able to reach ten draws in a row. This could bias the data, so in this experiment, I allowed participants to end an ongoing winning streak by either losing a game or drawing ten times after 50 games. This ensured that winning streaks could start until Game 50. Thus, the participant could potentially play 59 games. The most games played by a participant in this experiment were 56, ending a successful ten-game winning streak. Interaction with the experiment from the color vision test to the demographic questionnaire was done using the mouse on the computer. The participant used the right mouse button to advance the exper-

[3]For this test, I used five items from Velhagen and Broschmann (1995) that test for monochromacy and dichromacy

iment when prompted and the left mouse button to select items. Selected items were indicated by a black ring around the computer-selected item and a gray ring for the player's selection. Each game followed the same procedure. The computer opponent would select a circle, then the player. This would repeat until either the computer won by a winning triplet or all items were selected without the computer having a triplet. When the computer was able to create a triplet, the triplet was highlighted in red so that the player could see why he had lost, and the player could start the next game by clicking the right mouse button. If the game was tied, this was acknowledged, and the player could also advance with the right mouse button.

Materials

The tic-tac-toe isomorph was carefully designed to represent that of the original study by Zhang (1997) and to eliminate confounding factors. The experiment was conducted on a computer, the participant interacted with the experiment using the mouse. The experiment was performed in full-screen mode (screen size 19inch, resolution 1280 $\times$ 1024 pixels). Throughout the experiment, the background was uniformly white. All instructions and the game board were displayed on the screen. The experiments were implemented using Processing[4] a programming language for visual representations.

Computer opponent. The computer opponent used the same strategy as in Zhang (1997, Strategy 1, p. 215). For a summary of the strategy, see Table 5.5. This represents a perfect game strategy of several for tic-tac-toe. The computer opponent marked its choices with a black circle around the chosen element. Computer choices were always performed 500ms after the start of the game or the player's choice.

The tic-tac-toe isomorph. The color isomorph of tic-tac-toe was designed so that only the spatial arrangement differs between conditions and most confounding factors are randomized. The game board consisted of nine black circles, each with colored dots inside. There were 60 different board configurations for each condition, so each game had a different game board and none would repeat for each participant.

The isomorph was created by first representing the original tic-tac-toe in terms of its crucial components and their properties. The game board of tic-tac-toe consists of nine objects (in the original, these are the cells) and eight winning triplets, each consisting of three objects. Winning triplets in the original are the rows (3), columns (3), and

[4] www.processing.org

Table 5.5.: Tic-tac-toe computer strategy

Phase	Action
Step 1:	Computer(C) selects a random corner element (2-dot)
Step 2:	Player's (P) turn
Step 3:	*IF* P took a side element (3-dot), C takes the middle element (4-dot)
	ELSEIF P takes the opposite corner element, C takes another corner element
	ELSE C takes the opposite corner element
Step 4:	P's turn
Step 5:	*IF* C can win, select that option
	ELSEIF P can win next turn, C blocks P
	ELSE P started side, C takes corner element, which creates an impasse
Step 6:	P's turn
Step 7:	*IF* C can win, choose this option
	ELSE Block P from winning next turn
Step 8:	P's turn
Step 9:	C takes the remaining element

Note. For ease of communication, the description refers to the layout arranged in Condition 1. However, the same elements are selected regardless of the condition.

diagonals (2). In the color isomorph, the objects are circles with dots in them. The winning triplets are the three circles that have a dot of the same color.

The tic-tac-toe board also has symmetries such that all corner and side elements in the problem structure are identical. Selecting the top left corner first is equivalent to choosing any other corner first. The same applies to the side elements. In the isomorph, this is seen by the different number of dots within each circle.

Colors. The colored dot in each element was represented in one of three color schemes. In addition to the color schemes, I created twenty different distributions of colors for the winning triplets, resulting in 60 distinct color profiles. In this way, each element in each trial contained different colors, either by switching the color scheme or by color distribution. The colors used can be seen in Table 5.6.

All colors were taken from colorbrewer.[5] The website provides color schemes for displaying map data. These are tested for distinctness and are available for sequential as well as qualitative data. I used this resource to get easily distinguishable colors for the eight winning triplets. I chose three of their color schemes for qualitative data with eight classes, namely *set1*, *dark2*, and *paired*. Colors from one color scheme were used in

[5] www.colorbrewer2.org

Table 5.6.: Color schemes for tic-tac-toe

Color scheme	R	G	B
set1			
	228	26	28
	55	126	184
	77	175	74
	152	78	163
	255	127	0
	255	255	51
	166	86	40
	247	129	191
dark2			
	27	158	119
	217	95	2
	117	112	179
	231	41	138
	102	166	30
	230	171	2
	166	118	29
	102	102	102
paired			
	166	206	227
	31	120	180
	178	223	138
	51	160	44
	251	154	153
	227	26	28
	253	191	111
	255	127	0

Note. All colors used in Experiment 2. The three color schemes are from colorbrewer (www.colorbrewer2.org) for eight classes of qualitative data.

each game. To further randomize the colors per trial, I created twenty variations of the color distribution across the winning triplets. As there are eight winning triplets, there are 8! permutations of the colors over these triplets. For abstraction, I numbered the colors according to their position from one through eight, as shown in Table 5.6. Each permutation was then represented as 8-tuples of color identifiers, with each index representing one winning triplet. The possible permutations were generated using MATLAB, which provides an enumeration of the permutations. So there is a pattern in the output that risks regularities in the permutation, where one takes the naïve approach of taking every $8!/20 = 2016$th permutation. I therefore aimed to choose permutations that were as diverse as possible while sampling from the entire range of permutations. To do this, I analyzed samples of twenty permutations compared to a uniform distribution of colors. Each sample was analyzed for how often colors occurred at each index of the 8-tuple, i.e. how often the same color represented the same winning triplet. The optimal uniform distribution is the hypothetical distribution where each color occurs 2.5 times per winning triplet. Per sample, the count of each color occurring at each winning triplet was calculated, and then the RMSE of this is calculated compared to the uniform distribution of 2.5 per winning triplet. In this way, eight RMSE values, one for each winning triplet, were calculated. To compare the different samples, I used the sum of these eight values. I generated ten samples of twenty permutations. For comparison: The naïve implementation of each 2016th permutation was one sample. For the other samples, prime numbers (p_i) were chosen to reduce the probability of patterns emerging, and every p_i-th permutation was selected. If $8! \mod p_i > 20$, only the first twenty permutations were selected. The prime numbers chosen were: 2011, 2003, 1597, 1129, 947, 929, 691, 263, 211, plus the naïve 2016. The resulting sums of the RMSE were: 9.44, 9.38, 10.31, 9.25, 10.52, 10.03, 11.36, 16.32, 17.29, and 18.90, respectively. As it turned out, the higher primes give better results than the lower ones, and all primes are better than the naïve sample. As mentioned earlier, the permutations generated are an enumeration, and thus the naïve sample captures regularities in this enumeration, resulting in less randomness in the sample. The lower primes generate more candidates, and since I only selected the first twenty of these candidates, the sample is not representative of the population of permutations. The sample I chose for the experiment was based on $p_2 = 2003$. This sample performed well in terms of RMSE and also had no color that occurred more than four times in a winning triplet, while the best performing sample(p_1) had one color that occurred five times in one triplet.

The same color permutations were used for each color theme, so I was able to ensure 60 distinct color profiles. These were presented in random order per participant.

Element positioning per condition. Regardless of condition, all elements were presented within a virtual 400 $\times$ 400 pixel game area. The game area was indistinguishable from the background for the participant. Each element consisted of a black ring ($d = 80$px, stroke 2px) with 2-4 colored dots($d = 24$px). The three conditions of

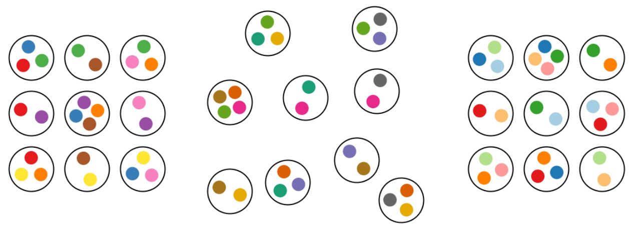

(a) C1 (*set1*) (b) C2 (*dark2*) (c) C3 (*paired*)

Figure 5.2.: Experiment 2: Example configurations of each condition with the respective color scheme in parenthesis.

this experiment differed only in the spatial arrangement of the elements of each game. The different arrangements are shown in Figure 5.2. In Condition 1 (C1), the elements in a grid were arranged such that the winning triplets corresponded to those of the classical tic-tac-toe(Fig. 5.2a). In Condition 2 (C2), the elements were arranged in such a way that the problem structure and the spatial structure of the game board did not match, and the spatial structure provided few grouping cues (Fig. 5.2b). And finally, in Condition 3 (C3), the spatial structure represented that of C1, but the problem structure was in conflict with the spatial structure (Fig. 5.2c).

In all conditions, each element was randomly rotated as it was drawn, so that the dots did not always in the same location within an element between games. Thus, randomization was performed by rotating the elements and the color profiles in all conditions. The positioning of the elements was done differently for each condition. In C1 & C3, the positions of the elements were arranged in a 3 by 3 grid. The center object was in the middle of the screen. The distance between adjacent elements was 20px, so the center of the elements was shifted by 100px. In C1, the winning triplets coincided with the rows, columns, and diagonals of the grid, just as in traditional tic-tac-toe. Randomization was performed only by the rotation and color profiles in C1. In C3, the problem structure should be in contradiction with the spatial structure, i.e. the elements had to be arranged so that the winning triplets do not coincide with rows, columns, and diagonals. Finally, the positioning in C2 should be random and should not have an unintentional correlation between spatial structure and problem structure.

In C2, the positioning of the elements should be such that the spatial structure does not correspond to the problem structure. This means that the overall structure should not be a grid or otherwise grouped to suggest that the elements are considered as a group. Moreover, the symmetry groups should not be grouped as such, i.e. not all 2-dot elements should, for example, be pooled together. The configurations for C2 were created by generating random configurations of elements and then discarding those that did not meet the C2 requirements. I generated 100 random configurations by distributing elements in the game area with a minimum distance between elements of 20px. The random configurations were then visually inspected to remove those configurations that exhibited accidental confounding grouping. First, I removed the configurations that did not seem to be evenly distributed across the game area. These were configurations that, for example, bundled all elements in one area, divided the elements into groups, or recreated a grid structure. In a second pass, the symmetry groups on each element were marked, and their distribution in each configuration was evaluated. All configurations that combined symmetry groups were removed. In this way, I was able to reduce the set of 100 configurations to 60 that had little to no obvious correlation between spatial structure and problem structure. Doing this via a visual inspection of the configurations may seem "unscientific", but as mentioned earlier, humans are extremely skilled at recognizing structure and regularity in visual configurations. This method allowed a quick yet reliable selection of configurations of elements for the experiment. An example of this is shown in Figure 5.2b(Appendix A.2 shows all random boards used). Configurations were presented in random order per participant in C2.

The goal for the configurations for C3 was that they would be as different as possible from C1 and that the winning triplets be represented as little as possible in the spatial structure. To achieve this, I started from the configuration of C1 and gave each element an index as if they were in a matrix. I then coded them in terms of their symmetry group and calculated all possible permutations of the symmetry groups over the 3 by 3 matrix. I have maintained 160 distinct permutations, i.e. which cannot be converted into each other by rotation or mirroring. To further reduce the set, all those permutations that had three 2-dot or 3-dot elements in a row, column, or diagonal were removed. The set was further reduced by all those permutations that had many adjacent 2-dot or 3-dot elements. I stuck with 38 permutations. For these, the possible distributions of the items were analyzed. I wanted to minimize the number of winning triplets while remaining consistent with the spatial structure. This was done exhaustively by applying all permutations of items over the generated symmetry group permutation and counting how many winning triplets were still in a row/column/diagonal, how many colors were in adjacent elements (horizontally, vertically, and diagonally), and if no adjacent colors were found, how many duplicate colors were in each row/column/diagonal. The result was that the least possible number of multiple colors in a row/column/diagonal was no triplet, seven neighbors, and five duplicate colors. Of these permutations, which differ as much as possible from C1, I have chosen eight. I created 60 configurations by applying rotation and mirror transformations to each, producing eight isomorphs per permutation. For a permutation, only the mirror and 180° rotation isomorphs were generated to obtain a total of 60 configurations(Appendix A.3 shows all distinct random boards before rotation and mirroring).

Practice trials. Before starting the actual games against the computer opponent, the participant had to complete a short exercise to familiarize him/herself with the elements and interaction with the experiment. The exercise consisted of two phases. In each phase, the participant had to select two out of three elements that contained the same color dot. Practice trials were repeated until the participant had successfully selected the correct item in successive trials. In Phase 1, the three elements were all 3-dot elements. Between phases, an instruction was presented on the screen, informing the participant that not all elements contained the same number of dots, and the exercise was repeated with a 2-dot, 3-dot, and 4-dot item. The prompt for the color was presented directly above the center of the screen and below the three elements of the trial. In C1 and C3, the elements were presented side by side, vertically aligned, with 20px horizontal space between items, while in C2, the three elements were additionally shifted vertically (-50px, 20px, -30px). Practice trials were presented in random order, and item positioning was also randomized between trials. All color schemes of the experiment were also used in the exercises. At the end of the exercise, the participant could start the experiment by himself.

Questionnaire. After the games were complete, the participant answered a short pen-and-paper questionnaire. The participants indicated gender, age, nationality, how they learned about the experiment, and the course in which they were enrolled in. They were also asked to describe their strategy, whether they knew tic-tac-toe, and whether knowing about it helped them. Finally, there was room for general comments on the experiment(Appendix A.2).

Procedure

The participants signed an informed consent form and received written instructions about the experiment. Thus, the aim of the study was to investigate *game-playing behavior*. Participants were informed of their task and that the goal was to tie the computer opponent and that there was a strategy they could use to consistently accomplish this. They were also told that each draw would earn them an additional €0.05, up to a maximum of €2.50. The experimenter then performed the color vision test. The experimenter asked the participants if everything was clear. After confirmation or clarification of any questions that arose, the participant was left alone in the lab to perform the experiment on the computer. The participant first completed the practice trials and then played 50 games against the computer opponent. Upon completion of the 50 trials, the computer screen displayed the number of ties achieved and the prompt to call the experimenter. The participant then completed the questionnaire and was compensated and thanked by the experimenter(Appendix A.1).

5.2.2. Results

The results of the tic-tac-toe experiment are summarized in Table 5.7

Table 5.7.: Results of the Tic-tac-toe experiment

	Consistent(C1)	Random(C2)	Conflicting(C3)
n	21	20	19
Strategy found	16	7	7
Proportion of sample	.73	.35	.39
		M(SD)	
Games until strategy	17.63(9.71)	25.00(9.56)	24.23(14.51)
Proportion of ties before strategy found (only those who found a strategy)	.23(.24)	.12(.14)	.22(.25)
Proportion of ties before strategy found (all participants)	.22(.22)	.09(.10)	.12(.17)
Mean reaction time of the first two moves before a strategy was found in seconds	10.50(8.13)	12.07(5.68)	9.58(4.82)

Note. Different results of the tic-tac-toe experiment.

"Games until strategy" are defined as the number of games until, and including, the first game of the win streak. At which point I assumed a successful strategy was formed. This statistic does not include participants who did not find a strategy over fifty games. Reaction times are the average time to make a choice in the first two moves of each game. The moves from the third option onward are only counter-moves if played correctly, and the player can be guaranteed to make two moves for each game, so the first two were analyzed. I analyzed the results in two steps. First, using multivariate analysis of variance (MANOVA) and then using posthoc analyses of variance (ANOVA). In the MANOVA, the independent variable was the condition and the dependent variables were the reaction time, the number of draws before a strategy was found, and the games until the strategy was found. Post-hoc analyses were conducted via ANOVA's measuring the effect of condition on response time, the tie proportion, and to test the start of the winning streak, respectively. MANOVA revealed a main effect of condition on the three performance measures (Wilks' lambda $= .73, F(2,57) = 3.19, p = .006$). ANOVAs showed a significant effect on the proportion of ties before the strategy was found ($F(2,57) = 3.19, p = .049$), a significant effect on the games until the winning streak ($F(2,57) = 5.96, p = .004$), and no significant effect on reaction time ($F(2,57) = .73, p = .49$).

5.2.3. Discussion

The results of the tic-tac-toe experiment show that the spatial structure of a problem representation is perceived and used by problem solvers when it is considered useful. It was easier to find a winning strategy for C1 than for C2 and C3. This is consistent with the results of Zhang (1997), who also found that the color version was more difficult than the traditional tic-tac-toe. Performance on trials before a strategy was found was also better for C1 than for the others when considering all participants, including those who never found a strategy. From those who eventually found a strategy, performance in the games before the strategy was nearly identical, which may indicate that the spatial structure is particularly important in finding the problem structure and the winning strategy. (Compare Experiment 3 of Zhang, 1997). The spatial structure of C3, even if contradictory to the problem structure, seems to have helped participants to organize the visual information and thus to facilitate the search for elements in the game. This is also reflected in the slightly higher reaction times in C3 compared to C2 – however, this difference was not statistically significant. Apart from this non-significant difference in reaction time, participants' performance in C2 and C3 was the same. Thus, the hypotheses that a conflicting spatial structure would inhibit performance are not supported with this experimental setup. It seems that spatial structure is sought as an aid but is quickly ignored as a structural aid when the solver found it to be of no help for this. Thus, the methodology was not appropriate to find such an effect. It could be that unwanted spatial structure effects such as the Simon effect (J. R. Simon, 1990) and spatial flanker tasks (Kopp et al., 1996) are detected at such an early stage of processing that this effect is not yet suppressed. Thus, the effect of incompatible spatial structure could be an attention-altering or priming effect that does not last long enough to show an effect in this task. The reaction times for the first two moves were each about ten seconds, much longer than the millisecond ranges in which these traditional S-R paradigm effects are measured.

5.3. Experiment 3: Spatial vs. Visual Features in Problem-Solving

The third experiment was the card sorting experiment. With this experiment, I wanted to investigate the relative importance of spatial and non-spatial information types. The empirical results were also compared with those of cognitive computer models performing the same task.

5.3.1. Method

Participants

The experiment was conducted online and ran in the participant's web browser. Participants were recruited via the platform "CrowdFlower"[6]. CrowdFlower specializes in outsourcing small, repetitive tasks like verifying contact information or tagging images by their users in exchange for monetary compensation. It is not primarily used as an experimental platform, unlike Amazon's Mechanical Turk, which was not available in my region. This means that the participants in this study did not have as much experience with psychological studies as many students in other studies. However, the game-like nature of the task nonetheless allowed even inexperienced participants to perform the task as instructed without difficulty. The platform also allows participants to be selected based on their estimated trustworthiness from the previously completed jobs. For this experiment, I selected only participants who had demonstrated at least a basic level of trustworthiness in ten previous jobs in order to have participants who

[6]www.crowdflower.com

followed instructions and were motivated to perform the task seriously. A total of 48 people participated. The participants were paid \$0.75 for their participation, and the task took about ten minutes. The amount represented an average hourly compensation of \$4.35, and participants' average satisfaction with the compensation was 4.5 out of 5 on their voluntary post-job survey ($N = 27$).

Three participant's answers were below the guessing probability, so it is assumed that they did not attempt to solve the task. Three participants also did not successfully find a sorting criterion, two of whom were also in the earlier group of three guessing participants. Data from all four were not analyzed further. This left data from 44 participants; 11 female, $M_{age} = 31.8$ years, $SD = 9.2$, age-range 18–55.

Task

The task was essentially that of the Wisconsin Card Sorting Test (E. A. Berg, 1948). Participants were provided with an online version of the task. In the upper part of the window, there were three target cards, and below them, the sorting cards were presented one after the other. Each card had geometric shapes printed on it with four different characteristics, each of which could be in three states. These characteristics were color, fill, topology, and location. The task for the participants was to assign each card to a target card by clicking on it. The participant was then given feedback on whether the card was sorted correctly, and the next card was presented. When the participant had correctly sorted ten cards in a row, the sorting criterion was changed unnoticed by the participant. The task ended when the participant correctly sorted ten consecutive cards for five sorting criteria, or each card was presented twice (156 trials). The order of the sorting criteria was predefined and always based on one of the four characteristics of the geometrical shapes. After the card sorting task, the participant completed a short demographic questionnaire.

Materials

The experiment was implemented in JavaScript and hosted on formsite.com[7], The demographic questionnaire was also conducted on this site.

The stimulus cards for this experiment had to meet a number of criteria. There are graphically represented objects on each card. These objects differ in four characteristics. Two characteristics include spatial and non-spatial information types, respectively.

[7] www.formsite.com

Each characteristic can be in different states, and each characteristic has the same number of states. The states of the characteristics can be varied independently between the cards. All characteristics should be presentable in one visual area so as to not to change modality between information types. The characteristics should clearly be only spatial or include non-spatial information types. The characteristic of orientation can, for example, be interpreted as spatial by applying a rotation operation between cards, but the same object can also be perceived as a different, new object. The shape as a characteristic can be spatial in the relations of the vortices of the objects or non-spatial in an object as a token perception. In addition, the use of shape as a characteristic gives rise to the problem that it is difficult to balance salience between shapes. Finally, I chose four characteristics to use on the cards. The two spatial characteristics were *Location* and *Topology*. The non-spatial ones were *Color* and *Fill*. These characteristics meet the requirements for the representation of an information type and can be communicated well via geometric shapes.

The cards were 170×255 pixels in size and were constructed as shown in Figure 4.5. A black rectangular outline with rounded corners was drawn, analogous to real playing cards. A rectangle and a circle were printed on these cards. Both shapes were vertically aligned, and the rectangle was larger than the diameter of the circle so that the circle could lie inside the rectangle without touching any edges. The topology characteristic was constructed by changing the horizontal positioning of the shapes. The three states were *the rectangle contains the circle*, *a partial overlap*, and *disconnected*. In the containment state, the circle and rectangle were both horizontal and vertical. In the partial overlap state, the center of the circle was horizontally aligned with the right edge of the rectangle. In the disconnected state, the circle was to the right of the rectangle. The location characteristic's states were top, middle, and bottom. For this purpose, the two forms were placed at 25%, 50%, or 75% of the height of the card. Three colors were chosen that could be distinguished by color-blind participants namely blue (rgb: 44,155,179), red (rgb: 179,8,77), and yellow (rgb: 230,210,32). The outline of the shapes, as well as the fill if any, were shown in this color. Finally, the fill can be either empty, striped, or filled. The fill in the filled and striped states was shown in the same color but with 50% opacity. This was done so that in the partial overlap state and in the containment states of the topology, it was apparent that there were two shapes, not one complex shape. In the striped state, diagonal white stripes were drawn so that white and the respective color formed stripes of equal thickness.

There was one card for each possible characteristic-state combination (81 cards), three of which were used as target cards. These were the cards that had each characteristic in

State 1, 2, and 3, respectively. The state indices are essentially arbitrary, and the target cards were shown side by side in a random order per participant. The remaining 78 cards were used as sorting cards(See Appendix B for all cards).

The target cards were displayed centered at the top of the browser window, with space for feedback text below, and sorting cards displayed below. The size and spacing between elements were chosen so that they were fully visible even on relatively small computer screens. The minimum size used was a 13 inch 1024 $\times$ 768 pixel display running Mac OS X 10.6.

The short post-test questionnaire was conducted on a second page in the same browser window, available after card sorting was completed.

The order of the cards was randomized for each participant. In order to give participants enough cards to potentially find all criterion changes, the same random order was repeated after the 78 cards had been seen, so that a total of potentially 156 sorting cards were presented.

The order of sorting criteria order was counterbalanced between the participants. Each switch meant that the new criterion could either be spatial (S) or non-spatial (N-S). This results in the following possible switches: S→S, S→N-S, N-S→N-S, N-S→S. Each participant should have to solve each of these switches, so four new sorting criteria plus the initial one. I have defined the sort order conditions using a circular linked list. The sorting criteria were presented as the following list: Color(N-S), Fill(N-S), Topology(S), Location(S) Each participant had to solve five criteria from this list, i.e. the first one would be repeated. I varied the initial item (4 possibilities) and the direction of the pass (2 possibilities) so that there were eight different sort order conditions.

Procedure

The participants received a referral link on the recruitment website (crowdflower.com) and a participant number they had to enter on the experiment website (formsite.com). This participant number was used to determine their condition and to initialize the card sorting task. Upon reaching the experiment's site, participants were presented with the following introductory text explaining the task.

> Welcome to the experiment. In this experiment, you are asked to sort cards to either one of the three target cards presented on the top. You sort the cards by clicking on the target card that you want to sort the presented card to. After each sorted card, you will get feedback on whether the card was

> sorted correctly. The goal is to find the criterion by which the cards should be sorted. Please enter your participant number below to start.

After providing the participant number and confirmation, the card sorting task was started, and the first sorting card was displayed. After the participant clicked on a target card, feedback was provided below the target cards. This was either "CORRECT" in green or "FALSE" in red. After a delay of 600ms, the next card was presented, during this delay, no input was possible to prevent accidental input. After ten consecutive correct trials, the sorting criterion was changed. This procedure was repeated until the participant either had ten consecutive correct trials for each sorting criterion or had seen each sorting card twice. The participant was then asked to follow a link to a short questionnaire. The questionnaire asked for age, gender, nationality, and a brief description of the strategy used.

After submitting the questionnaire, the participant was given a response code. Providing this code on the CrowdFlower website completed the task, and the participant received his or her compensation.

5.3.2. Information Type Salience Model

For the Information Type Salience Model (ITSM), a version of the task was implemented that behaved identically to the experimental task but could be interacted with by the computer model. The ITSM was developed with the goal of first creating a very basic version and then study its behavior. Changes would be made to reflect theories of human strategies for WCST. Finally, the behavior of the model versions would be compared to the performance of the human participants. Similar performance was an indication that the processes in the computational model mirrored those of the participants.

The ITSM had internal representations for its hypothetical sort criterion and saliency for various sorting criteria. It could be presented with a card and feedback of a sorting attempt. After receiving feedback from the previous trial and the card for the current trial, the model outputs the target card to which it wants to sort the card. While the model received positive feedback, it stayed with the same hypothetical criterion. Behavior upon receiving negative feedback differed between model versions depending on which salience it represented.

The initial basic version (V1) of the model represented salience for the four sorting criteria and, when receiving negative feedback, selected one of the four depending on its salience. It could continue with the same criterion for which it had just received negative feedback. With flat salience values for the four criteria, the model performs extremely well and well above the level of the human participants. The goal of this experiment was to determine whether spatial information is simply more salient than non-spatial information. In order to estimate the true relative salience values of the four represented information types, the salience values were determined dynamically. Each model version performed the task, and its results were compared with the empirical data. Saliency was adjusted to increase the fit of model results to empirical results. Model fit was determined by measuring the root mean squared error (RMSE) between the result vectors. These vectors contained the data from Table 5.8 for the empirical vector and the model data for the model. Saliencies were optimized by using a Metropolis algorithm to increase the fit. After optimizing the saliencies and minimizing the RMSE between the model and empirical results, the resulting salience values predict the true salience values. In addition, model consistency was tested.

To optimize the salience, the model performed an experiment run as a participant would. The results of this run were compared with the empirical results by calculating the fit (RMSE). Then the saliencies were optimized and the procedure repeated.

The Metropolis algorithm works by generating new values for saliency based on a probability distribution. For each iteration of the algorithm, the saliencies are changed, and if the fit improves, the new saliencies are used to start the probability distribution. After a certain number of iterations, the standard deviation(SD) of the probabilities is reduced to obtain a finer and finer determination of the optimized saliency. The algorithm has the following parameters: the number of iterations or jumps performs, the probability distribution, the SD of the probability distribution, the step-down intervals at which the SD is reduced, the amount by which the SD is reduced at each step, and the measure used to determine the fit.

The following Metropolis parameters were used. Normal distribution was used for each jump. The starting SD was 0.25 while starting saliency values were randomly chosen from the range of 0-1, excluding 1. The iteration was set to 20000 with SD reduced every 200 jumps, which reduced the SD by a factor of 0.985. The fit was measured by the root mean squared error (RMSE) between the modeled and empirical results.

Each modeled experiment was run between 5 and 50 times with 44 “participants” each. Results of these runs were averaged, and the RMSE of these averaged results to the empirical results was used as a measure of fit.

The different versions of the model differed in what criteria they considered and whether they could stick to the criterion they had just received negative feedback. I will discuss the following models in this analysis.

1. 4 saliencies for each graphical category, the switches are determined only by the probability distribution. A change to the previous hypothesis is possible.
2. 4 saliencies for each graphical category, the switches are determined only by the probability distribution. It is not possible to switch to the previous hypothesis.
3. 4 saliencies for graphical categories, 2 factors for graphic type categories visual and spatial. Factors multiply the saliency in the category.
4. 5 saliencies, 4 graphic categories 1 for best fit. The best fit chooses the target with the most categories corresponding to one target. If two targets each match two categories, the model selects one at random.
5. 4 saliencies, 2 factors for switching between information types: One factor for a switch from spatial category to visual category and one factor for switches from visual to visual. Visual to spatial and spatial to spatial were the reciprocal of the factor, respectively. These factors were used to determine salience depending on which previous hypothesis the model contained.

5.3.3. Results

Empirical study

For each type of switch and the four starting criteria, the number of trials until the ten-card streak was reached and were recorded. The full results can be seen in Table 5.8. The lower the mean number of tries, the faster the participant found this criterion in the task. Assuming a naïve approach, the player first tries the most salient type of information. Thus, a high salience results in a low mean number of tries until the criterion is determined. At first glance, it is clear that the participants needed a considerable number of trials for each switch. A strategy of systematically testing the four criteria would require, at worst, four trials plus possible false positives in guessing. The comparison of the starting criteria shows that sorting by topology was most easily recognized by the participants($M = 5.58$). Fill came second and color and location were the least well recognized.

The switches from one criterion to another show no obvious evidence of an interaction of spatial versus non-spatial information types. It is not evident that switching from a spatial information type to a non-spatial one or vice versa is more difficult than switching within a category.

The analysis of these results was further carried out using the computational model.

Table 5.8.: Empirical results of card sorting

Criterion switch	n	M tries	SD tries
Starting with Color	12	20.75	22.59
Starting with Fill	11	12.91	9.87
Starting with Topology	12	5.58	8.83
Starting with Location	9	27.33	30.49
Color to Fill	16	11.13	8.39
Color to Location	10	16.00	13.13
Fill to Color	14	15.29	22.06
Fill to Topology	19	11.47	13.38
Topology to Fill	13	11.69	11.41
Topology to Location	14	38.71	28.18
Location to Color	14	6.86	6.55
Location to Topology	8	23.25	26.44

Note. This table shows, for each sort criterion switch, the number of successful observations, the mean trials until the onset of the ten-card-correct streak, and the standard deviation of that number.

Computational model

To assess the performance of the models, their respective consistency in optimizing the saliency values and their consistency in running the experiments with the same saliency values were recorded. The performance of Version 3 of the model was so erratic and inconsistent that further analysis was not possible. Model optimization generally did not improve after the first half of Metropolis iterations, indicating non-replicable results. Only models 1, 2, 4, and 5 are discussed in the results.

The consistency of the optimization was compared by repeating the optimization algorithm fifteen times with random saliency values given. The average achieved fit and its standard variance were used to compare optimization versions that differ in the number of experiments performed per Metropolis jump. In order to obtain reliable results, it is important that the Metropolis algorithm does not simply get stuck at saliency values that were achieved by a lucky trial run. Therefore, for each jump, the average of several modeled experimental runs was used to calculate the fit. Table 5.9 summarizes the results of different numbers of runs per Metropolis iteration. As can be seen, the variance between optimizations does not decrease uniformly. The achieved fit between optimization runs is relatively constant.

The consistency of the model version was tested by running the model fifteen times

Table 5.9.: Model optimization consistency results

RPI	Mean RMSE of Optimization	Mean Metropolis optimized parameter values(*SD*)					
		Model version 1					
		Color	Fill	Topology	Location		
5	6.78(0.38)	0.074(0.009)	0.706(0.244)	0.183(0.243)	0.037(0.007)		
10	6.83(0.10)	0.077(0.012)	0.801(0.011)	0.085(0.008)	0.037(0.005)		
15	7.30(0.56)	0.118(0.172)	0.575(0.321)	0.269(0.302)	0.038(0.004)		
20	7.01(0.05)	0.069(0.007)	0.803(0.013)	0.089(0.009)	0.039(0.004)		
50	7.20(0.26)	0.075(0.009)	0.751(0.178)	0.135(0.180)	0.039(0.003)		
		Model version 2					
		Color	Fill	Topology	Location		
5	8.84(0.17)	0.197(0.227)	0.463(0.160)	0.319(0.251)	0.021(0.004)		
10	8.98(0.19)	0.198(0.202)	0.512(0.103)	0.269(0.209)	0.020(0.003)		
15	9.08(0.19)	0.246(0.233)	0.523(0.113)	0.208(0.161)	0.023(0.004)		
20	9.09(0.20)	0.194(0.201)	0.479(0.109)	0.305(0.232)	0.022(0.004)		
50	9.10(0.15)	0.140(0.209)	0.527(0.119)	0.312(0.159)	0.021(0.003)		
		Model version 4					
		Color	Fill	Topology	Location	Best fit	
5	5.03(0.19)	0.101(0.018)	0.174(0.031)	0.115(0.016)	0.052(0.007)	0.558(0.035)	
10	5.30(0.14)	0.100(0.011)	0.177(0.040)	0.119(0.015)	0.050(0.004)	0.553(0.034)	
15	5.37(0.09)	0.094(0.010)	0.154(0.031)	0.114(0.014)	0.049(0.005)	0.588(0.025)	
20	5.45(0.11)	0.095(0.009)	0.163(0.033)	0.119(0.011)	0.051(0.005)	0.572(0.031)	
50	5.63(0.06)	0.093(0.007)	0.157(0.024)	0.118(0.010)	0.050(0.005)	0.582(0.023)	
		Model version 5					
		Color	Fill	Topology	Location	S to V factor	V to V factor
5	6.87(0.55)	0.090(0.094)	0.459(0.143)	0.333(0.087)	0.118(0.053)	0.634(0.292)	0.806(0.188)
10	6.98(0.44)	0.072(0.076)	0.486(0.159)	0.319(0.093)	0.122(0.049)	0.761(0.204)	0.830(0.076)
15	6.88(0.08)	0.045(0.010)	0.509(0.103)	0.315(0.083)	0.131(0.036)	0.858(0.110)	0.830(0.074)
20	7.17(0.48)	0.084(0.102)	0.494(0.165)	0.307(0.105)	0.115(0.044)	0.753(0.229)	0.774(0.207)
50	7.35(0.54)	0.122(0.175)	0.451(0.173)	0.305(0.076)	0.122(0.040)	0.791(0.224)	0.806(0.089)

Note. Consistency checks for the optimization of the different model versions. Each optimization was run fifteen times with random starting saliences/parameters . RPI are the number of modeled experiments per Metropolis jump. The Mean RMSE of optimization is the mean fit of the fifteen reruns and the *SD*.

with the same saliency values and comparing the results. The saliency values given were the result of a preliminary optimization run. Different numbers of experimental repetitions per Metropolis jump were compared per version. The results are summarized in Table 5.10. Again, the results of the reruns were averaged and the standard deviation compared. In this case, the results indicate that the consistency of the models improved when at least ten reruns were performed per Metropolis repetition.

Overall the best-fitting results were achieved by Model Version 4. This version had an additional "best fit" sorting strategy that selected the target card that had the most characteristics in common with the sorting card.

5.3.4. Discussion

It is important to remember that the best fit does not mean that the performance on the task was also the best. In fact, the simple Models 1 and especially 2 performed significantly better at the task than the human participants. Problem-solving at the level of information types is possible and actually superhuman in this task. This suggests that the human participant did something other than using the different types of information separately. The fit of Model Version 4 suggests that the participant used a more holistic approach and sought targets that fit the most categories. This could also be a way to minimize mistakes, as this has the greatest chance of getting positive feedback. Moreover, it could be a case of confirmation bias rather than a strategy to exclude and determine the correct sorting criterion. The models had the advantage over the human participants in knowing that there were four categories that could differ. The models that worked exclusively with these categories then outperformed the humans and had a worse fit than the version that considered the holistic category of most similarities.

Table 5.10.: Model consistency results

RPI	RMSE optimization	Mean RMSE reruns(*SD*)	Metropolis optimized parameter values					
			Model Version 1					
			Color	Fill	Topology	Location		
5	6.56	7.61(0.36)	0.08	0.80	0.09	0.04		
10	8.41	8.83(0.17)	0.75	0.12	0.09	0.05		
15	6.99	7.56(0.10)	0.08	0.77	0.11	0.04		
20	6.99	7.62(0.21)	0.08	0.81	0.08	0.03		
50	7.16	7.52(0.09)	0.08	0.79	0.09	0.04		
			Model Version 2					
			Color	Fill	Topology	Location		
5	9.13	9.91(0.29)	0.55	0.35	0.07	0.03		
10	8.84	9.37(0.14)	0.05	0.50	0.43	0.03		
15	8.95	9.38(0.16)	0.04	0.33	0.62	0.02		
20	8.86	9.22(0.16)	0.03	0.70	0.25	0.02		
50	9.07	9.26(0.06)	0.04	0.35	0.60	0.02		
			Model Version 4					
			Color	Fill	Topology	Location	Best fit	
5	4.90	6.58(0.44)	0.10	0.16	0.10	0.04	0.61	
10	5.47	6.50(0.27)	0.11	0.24	0.10	0.05	0.50	
15	5.36	6.20(0.24)	0.11	0.13	0.13	0.05	0.59	
20	5.48	6.08(0.18)	0.10	0.15	0.12	0.05	0.59	
50	5.65	6.15(0.11)	0.09	0.13	0.11	0.06	0.61	
			Model Version 5					
			Color	Fill	Topology	Location	S to V factor	V to V factor
5	6.70	7.87(0.50)	0.04	0.30	0.46	0.20	0.77	0.94
10	6.84	7.52(0.17)	0.06	0.58	0.25	0.11	0.79	0.78
15	6.79	7.55(0.19)	0.03	0.33	0.47	0.18	0.92	0.93
20	6.88	7.54(0.22)	0.05	0.56	0.27	0.11	0.84	0.81
50	7.11	7.36(0.11)	0.07	0.76	0.12	0.05	0.82	0.56

Note. Consistency checks for the different model versions. Each version was run fifteen times with the saliency/parameters generated by the Metropolis optimization. RPI is the number of modeled experiments per Metropolis jump. The second column is the fit of the winning result of the optimization. The third column is the mean fit of the fifteen reruns with the optimized saliency/parameters.

6. Conclusion & Outlook

In this thesis, I described a theoretical framework for analyzing information use during problem-solving. I describe the framework of analyses of the level of abstraction of information types. I have defined several challenges for experimental methods to apply this framework. I then conducted pilot studies for the proposed experimental paradigms.

In this chapter, I will discuss the contributions of the work presented in this thesis to the state of the art, future work based on the results of this thesis, and the conclusions for the proposed framework.

6.1. Contributions

I propose a framework for analyzing human information use. The focus of cognitive research of problem-solving has been on object-level descriptions since these are entities that people deal with. For the subdomain of computational modeling of human cognition, one challenge is that for models operating in real-world scenarios, information sources and the type of information used must be explicitly specified. Even when people think about a problem at an object abstraction level, the information they gather during the problem-solving process differs. My framework allows problem-solving to be explored at this level of abstraction, facilitating the creation of accurate computational models. Much information is processed subconsciously because we are exposed to it throughout our lives. Similar to the perceptual-cognitive universals of Shepard (1994), information of certain types allows for certain conclusions and implies certain interactions with the world. Just as an object cannot teleport, perceiving it at two locations infers movement. The subconscious nature of these effects prevents the application of introspective methodologies of cognitive science normally used for problem-solving research.

In this thesis, I designed and piloted novel experimental paradigms to investigate the use of information types in problem-solving at an information type level of abstraction. These approaches allow to specifically test the influences of controlled presentation information types on the problem-solving process. These paradigms enable cognitive scientists interested in problem-solving behavior to manipulate the level of abstraction of information types in their studies (Schultheis & Wienemann, 2019, September 2–4).

The special role of spatial information for animated behavior and goal-directed action makes it a particularly suitable object of investigation in these pilot studies, the first new experiment aimed to test the transfer of information types between different domains. The use of distance information in a non-spatial domain was investigated. To accomplish this, I developed an analogy-building task. The task uses six term analogies that allow building a frame of reference within the analogy, thus representing relative distance information. The task proved suitable for manipulating distance information so that it could be transferred to a non-spatial target domain. This approach allows the study of effects at the abstraction level of information types in a domain closely related to human learning processes.

The second experimental paradigm was designed to test the effects of spatial information types that are not task-relevant, namely, relative orientation and distance. It has been found that the way problems are represented affects their difficulty for humans by changing the processes involved in representing the relevant information (Kotovsky et al., 1985). Spatial structure is a powerful enabler for perceiving problem structure. To apply my framework of information type abstraction level analyses, I adapted a tic-tac-toe game paradigm. In this version of the game, the spatial arrangement of the *moves* is irrelevant to problem-solving. The studied variable of the spatial arrangement is further hidden within the paradigm to allow the study of the emerging effects. This allows to manipulate the spatial structure of the problem representation and to study its effect on the solution behavior. I could detect a facilitating effect of the spatial structure but not an inhibiting effect of the contradictory structure. This result further supports the notion that human problem solvers seek spatial information even when it is not necessary for structuring their surroundings. However, the structure is also not blindly applied when it is not helpful. Rather I found small indications that in the condition set to interfere with problem-solving, participants used spatial structure to aid search, resulting in slightly faster play than a random configuration. These findings support the important role of spatial information for human problem-solving behavior.

The last paradigm I developed to apply the level of abstraction approach to information types aimed to test the use of competing information types. This was done to test whether spatial information is simply more salient than non-spatial information. Furthermore, subsequent analysis using a cognitive computer model allowed us to test whether a problem-solving approach at an information-type level of abstraction was viable. To accomplish this, I developed a card-sorting paradigm that allowed cards to be sorted by various spatial and non-spatial categories. The computer model I created was able to easily accomplish the task, suggesting the possibility of such an approach to problem-solving. It further suggests that human problem solvers do not employ such a strategy but rather shuffle information types into a more holistic category of fit.

The paradigm of this experiment allowed us to test the relative importance of different information types. Analysis via a computer model allows conclusions to be drawn about the strategy used by the participants.

6.2. Future Research

The experimental paradigms developed in this thesis have proven to be a viable way to study the information used in problem-solving tasks. Nonetheless, there are still a few kinks to work out with further use. I leave the work on optimizing the stimuli materials or experimental setups for future work.

Cross-domain analogy building could also be extended to other questions. The approach is novel in that it uses six-term analogies to create its own frame of reference in the base part of the analogy. I have not seen this approach in analogy research before, and I think there are more applications for it than just information-type use research. Further work on the use of different information types would also be very interesting. The paradigm could also be carried out without the explicit basis of analogy in the three-term series problem. This was made clear in this version by the arrangement of the letters on the screen. This was done because it was not known whether the participants were even able to solve the task without being able to offload information into the environment. However, the participants solved the task readily, and a more implicit mental model might influence the results.

Furthermore, it would be interesting to alternate the spatial structure in the tic-tac-toe experiment, for example, to further determine what structures are necessary to be recognized and used as helpful. Moreover, by gradually introducing randomness in the placement of the elements in this experiment, one could test how far deviations from the clear spatial structure are tolerated. In their study, Zhang (1997) found that spatial symmetry can help convey symmetry in the problem structure of tic-tac-toe. A variation of my paradigm with a pseudo-random placement of items so that spatial symmetry representing the problem structure might provide more insight into where a spatial structure is used and helpful. The fuzziness in the placement of the items could be varied to test for correlations of varying strength between spatial and problem structure. That said, the symmetries in the role of each item are present even in the current implementation by the number of colored dots in each circle. The spatial structure, as in many other examples, seems to be particularly suitable for conveying abstract structures. Spatially arranging all circles with the same number of dots together could be another option to test if the symmetry between the items can be communicated, and is worth a follow-up study.

The information-type level of abstraction approach to problem-solving should be implemented in other computational cognitive modeling systems. This would allow a broader application of the approach and subsequently lead to more insight into its explanatory power. The approach I have taken with the information type salience model can of course be extended to include more than just saliences. I think that computational understanding of what information types are being used in certain problem domains will enable more cognitively appropriate cross-domain systems of human-computer interaction. Understanding the types of information used to understand problem representations can lead to better representation of information in crucial systems where error reduction is a top priority.

6.3. Conclusions

The goal of this thesis was to test the assumption that problem-solving behavior can be studied at an information-type level of abstraction. I have provided the theoretical framework necessary to study this topic and developed tools to study information used at an information-type level of abstraction in cognitive experiments.

I maintain that the frameworks of a world structure based on relationships described by different types of information is a useful theory to describe the environment in which problem solvers operate.

I set out to explore this theory by addressing four challenges to this approach. Could I investigate the transfer of information types at an information-type level of abstraction, generate a task that allows manipulation of spatial structure and measure its spontaneous use by problem solvers, determine if spatial information was simply more salient than other information types, and create a computational model to exemplify problem-solving at an information-type level of abstraction? I have achieved examples for each of these challenges. I have discussed some limitations or shortcomings in the previous section, yet I maintain that this is thesis sets the stage for further exploration of information-type level of abstraction analyses of human cognition.

I offer the opportunity to examine the impact invariants in our world have on how our perceptions of this world are formed. The structure of the environment in which an organism functions determines the rules by which it must play. Being able to manipulate these elementary parts of our world structure and measure the effects this has on problem-solvers offers new possibilities for the cognitive science community.

Bibliography

Adriaans, P. (2012). Information. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Fall 2013 ed.). http://plato.stanford.edu/archives/fall2013/entries/information/

Ali, N., & Peebles, D. (2012). The effect of gestalt laws of perceptual organization on the comprehension of three-variable bar and line graphs. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *55*(1), 183–203. https://doi.org/10.1177/0018720812452592

Alibali, M. W. (2005). Gesture in spatial cognition: Expressing, communicating, and thinking about spatial information. *Spatial Cognition & Computation*, *5*(4), 307–331. https://doi.org/10.1207/s15427633scc0504_2

Allen, G. L. (1999). Cognitive abilities in the service of wayfinding: A functional approach. *The Professional Geographer*, *51*(4), 555–561.

Allen, G. L. (2000). Principles and practices for communicating route knowledge. *Applied Cognitive Psychology*, *14*(4), 333–359.

Anderson, B. L. (2003). Perceptual organization and white's illusion. *Perception*, *32*(3), 269–284. https://doi.org/10.1068/p3216

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), 1036–1060. https://doi.org/10.1037/0033-295x.111.4.1036

Anderson, J. R., Reder, L. M., & Simon, H. A. (1996). Situated learning and education. *Educational Researcher*, *25*(4), 5–11. https://doi.org/10.3102/0013189X025004005

Archer, A. F. (1999). A modern treatment of the 15 puzzle. *The American Mathematical Monthly*, *106*(9), 793–799. https://doi.org/10.2307/2589612

Ashcraft, M. H. (2005). *Cognition* (4th International edition). Pearson Education.

Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology*, *59*(1), 617–645. https://doi.org/10.1146/annurev.psych.59.103006.093639

Barsalou, L. W. (2009). Simulation, situated conceptualization, and prediction. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *364*(1521), 1281–1289. https://doi.org/10.1098/rstb.2008.0319

Berg, E. A. (1948). A simple objective technique for measuring flexibility in thinking. *Journal of General Psychology*, *39*, 15–22.

Berg, W. K., & Byrd, D. L. (2002). The tower of london spatial problem-solving task: Enhancing clinical and research implementation. *Journal of Clinical and Experimental Neuropsychology*, *24*(5), 586–604. https://doi.org/10.1076/jcen.24.5.586.1006

Bonifazi, S., Farnè, a., Rinaldesi, L., & Làdavas, E. (2007). Dynamic size-change of peri-hand space through tool-use: Spatial extension or shift of the multi-sensory area. *Journal of Neuropsychology*, *1*(1), 101–114. https://doi.org/10.1348/174866407X180846

Brainard, D. H., & Hurlbert, A. C. (2015). Colour vision: Understanding #TheDress. *Current Biology*, *25*(13), R551–R554. https://doi.org/10.1016/j.cub.2015.05.020

Bregman, A., & Campbell, J. (1971). Primary auditory stream segregation and perception of order in rapid sequences of tones. *Journal of Experimental Psychology*, *89*(2), 244–249. https://doi.org/10.1037/h0031163

Brown, J. S., Collins, A., & Duguid, P. (1989). Situated cognition and the culture of learning. *Educational Researcher*, *18*(1), 32–42. https://doi.org/10.3102/0013189x018001032

Brown, W. N. (1964). The Indian games of pachisi, chaupar, and chausar. *Expedition*, *6*(3), 32–35.

Bruner, J. S., & Minturn, A. L. (1955). Perceptual identification and perceptual organization. *Journal of General Psychology*, *53*, 21–28.

Brunyé, T. T., Andonova, E., Meneghetti, C., Noordzij, M. L., Pazzaglia, F., Wienemann, R., Mahoney, C. R., & Taylor, H. A. (2012). Planning routes around the world: International evidence for southern route preferences. *Journal of Environmental Psychology*, *32*(4), 297–304. https://doi.org/10.1016/j.jenvp.2012.05.003

Brunyé, T. T., Howe, J. L., & Mahoney, C. R. (2013). Acute exercise suppresses judgments of facial emotion intensity. *Motivation and Emotion*, *37*(4), 787–798.

Brunyé, T. T., Mahoney, C. R., Gardony, A. L., & Taylor, H. A. (2010). North is up(hill): Route planning heuristics in real-world environments. *Memory & Cognition*, *38*(6), 700–712. https://doi.org/10.3758/mc.38.6.700

Byrne, R. M., & Johnson-Laird, P. N. (1989). Spatial reasoning. *Journal of Memory and Language*, *28*(5), 564–575. https://doi.org/10.1016/0749-596X(89)90013-2

Campbell, M., Hoane, A., & Hsu, F.-h. (2002). Deep blue. *Artificial Intelligence*, *134*(1), 57–83. https://doi.org/10.1016/S0004-3702(01)00129-1

Campos, M. (2009). The notion of information. *tripleC*, *7*(2), 323–326.

Capurro, R., & Hjørland, B. (2003). The concept of information. *Annual Review of Information Science and Technology*, *37*, 343–411. https://doi.org/10.1002/aris.1440370109

Casasanto, D., & Boroditsky, L. (2008). Time in the mind: Using space to think about time. *Cognition*, *106*(2), 579–593. https://doi.org/10.1016/j.cognition.2007.03.004

Charness, N. (1992). The impact of chess research on cognitive science. *Psychological Research*, *54*(1), 4–9. https://doi.org/10.1007/BF01359217

Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, *4*(1), 55–81. https://doi.org/10.1016/0010-0285(73)90004-2

Chater, N., & Brown, G. D. A. (2008). From universal laws of cognition to specific cognitive models. *Cognitive Science*, *32*(1), 36–67. https://doi.org/10.1080/03640210701801941

Clancey, W. J. (2008). Scientific antecedents of situated cognition. In P. Robbins & M. Aydede (Eds.), *Cambridge handbook of situated cognition* (pp. 11–34). Cambridge, Cambridge University Press.

Clarey, C. (2014, February 22). Olympians use imagery as mental training. *The New York Times*. http://www.nytimes.com/

Conroy Dalton, R. (2003). The secret is to follow your nose. *Environment and Behavior*, *35*(1), 107–131. https://doi.org/10.1177/0013916502238867

Coughlan, G., Coutrot, A., Khondoker, M., Minihane, A.-M., Spiers, H., & Hornberger, M. (2019). Toward personalized cognitive diagnostics of at-genetic-risk alzheimer's disease. *Proceedings of the National Academy of Sciences*, *116*(19), 9285–9292. https://doi.org/10.1073/pnas.1901600116

Cox, R., & Brna, P. (1995). Supporting the use of external representations in problem solving: The need for flexible learning environments. *International Journal of Artificial Intelligence in Education*, *6*(2-3), 239–302.

de Bot, K., Lowie, W., & Verspoor, M. (2007). A dynamic systems theory approach to second language acquisition. *Bilingualism: Language and Cognition*, *10*, 7–21. https://doi.org/10.1017/S1366728906002732

Dehaene, S. (2001). Précis of the number sense. *Mind and Language*, *16*(1), 16–36. https://doi.org/10.1111/1468-0017.00154

Demakis, G. J. (2003). A meta-analytic review of the sensitivity of the wisconsin card sorting test to frontal and lateralized frontal brain damage. *Neuropsychology*, *17*(2), 255–264. https://doi.org/10.1037/0894-4105.17.2.255

de Voogt, A. (2003). Hawalis in oman: A first account of expertise and dispersal of four-row mancala in the middle east. *Board Game Studies*, *6*, 95–98.

Dewdney, A. K. (1988). *The armchair universe: An exploration of computer worlds*. Freeman.

Donders, F. C. (1869). Over de snelheid van psychische processen [On the speed of mental processes]. *Nederlandsch Archief voor Genees- en Natuurkunde*, *4*, 117–145.

Dreyfus, H., & Dreyfus, S. (1992). What artificial experts can and cannot do. *AI & SOCIETY*, *6*(1), 18–26. https://doi.org/10.1007/BF02472766

Eliasmith, C. (1996). The third contender: A critical examination of the dynamicist theory of cognition. *Philosophical Psychology*, *9*(4), 441–463. https://doi.org/10.1080/09515089608573194

Epstein, S. L., Gelfand, J., & Lock, E. (1998). Learning game-specific spatially-oriented heuristics. *Constraints*, *3*(2), 239–253. https://doi.org/10.1023/A:1009734028965

Ericsson, K. A., & Simon, H. A. (1980). Verbal reports as data. *Psychological Review*, *87*(3), 215–251. https://doi.org/10.1037/0033-295x.87.3.215

Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*, *16*(1), 143–149. https://doi.org/10.3758/BF03203267

Finkel, I. L. (2007). On the rules for the royal game of Ur. In I. L. Finkel (Ed.), *Ancient board games in perspective* (pp. 16–32). British Museum Press.

Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for "top-down" effects. *Behavioral and Brain Sciences*, *39*, 1–77. https://doi.org/10.1017/S0140525X15000965

Freksa, C. (1991). Conceptual neighborhood and its role in temporal and spatial reasoning (M. Singh & L. Travé-Massuyès, Eds.). In M. Singh & L. Travé-Massuyès (Eds.), *Procedings of the imacs workshop on decision support systems and qualitative reasoning*, Amsterdam, The Netherlands.

Freksa, C. (2015). Strong spatial cognition (I. S. Fabrikant, M. Raubal, M. Bertolotto, C. Davies, S. Freundschuh, & S. Bell, Eds.). In I. S. Fabrikant, M. Raubal, M. Bertolotto, C. Davies, S. Freundschuh, & S. Bell (Eds.), *Spatial information theory: 12th international conference, COSIT 2015, Santa Fe, NM, USA, October 12-16, 2015, proceedings*, Springer International Publishing. https://doi.org/10.1007/978-3-319-23374-1_4

Fu, W.-T., Hills, T., & Todd, P. M. (2015). Interfacing mind and environment: The central role of search in cognition. *Topics in Cognitive Science*, *7*(3), 384–390. https://doi.org/10.1111/tops.12153

Funt, B. V. (1980). Problem-solving with diagrammatic representations. *artificial intelligence*, *13*(3), 201–230. https://doi.org/10.1016/0004-3702(80)90002-8

Gallagher, S. (2008). Philosophical antecedents of situated cognition. In P. Robbins & M. Aydede (Eds.), *Cambridge handbook of situated cognition* (pp. 35–51). Cambridge, Cambridge University Press.

Garland, M., & Heckbert, P. S. (1997). Surface simplification using quadric error metrics, In *Proceedings of the 24th annual conference on Computer Graphics and Interactive Techniques*, New York, NY, ACM Press/Addison-Wesley Publishing Co. https://doi.org/10.1145/258734.258849

Gentner, D. (2001). Spatial metaphors in temporal reasoning. In M. Gattis (Ed.), *Spatial schemas and abstract thought* (pp. 203–222). Cambridge, MA, US, MIT Press.

Gentner, D., & Smith, L. A. (2013). Analogical learning and reasoning. In D. Reisberg (Ed.), *The oxford handbook of cognitive psychology* (pp. 668–681). New York, NY, Oxford University Press.

Gibson, J. J. (1986). *The ecological approach to visual perception*. Hillsday, NJ, Lawrence Erlbaum Associates. (Original work published 1979)
https://books.google.de/books?id=8BSLBQAAQBAJ

Giudice, N. A., Klatzky, R. L., & Loomis, J. M. (2009). Evidence for amodal representations after bimodal learning: Integration of haptic-visual layouts into a common spatial image. *Spatial Cognition & Computation*, *9*(4), 287–304. https://doi.org/10.1080/13875860903305664

Glenberg, A. M., & Kaschak, M. P. (2002). Grounding language in action. *Psychonomic Bulletin & Review*, *9*(3), 558–565. https://doi.org/10.3758/BF03196313

Gobbetti, E., & Marton, F. (2005). Far voxels: A multiresolution framework for interactive rendering of huge complex 3D models on commodity graphics platforms. *ACM Transactions on Graphics*, *24*(3), 878–885. https://doi.org/10.1145/1073204.1073277

Gobet, F., Retschitzki, J., & de Voogt, A. (2004). *Moves in mind: The psychology of board games*. Taylor & Francis https://books.google.de/books?id=CLt5AgAAQBAJ.

Goldman, J.-P. (2011). EasyAlign: An automatic phonetic alignment tool under praat (P. Cosi, R. D. Mori, G. D. Fabbrizio, & R. Pieraccini, Eds.). In P. Cosi, R. D. Mori, G. D. Fabbrizio, & R. Pieraccini (Eds.), *Interspeech-2011*.

Goodman, S. G., Seymour, T. L., & Anderson, B. R. (2016). Achieving the performance benefits of hands-on experience when using digital devices: A representational approach. *Computers in Human Behavior*, *59*, 58–66. https://doi.org/10.1016/j.chb.2016.01.006

Greeno, J. G., & Moore, J. L. (1993). Situativity and symbols: Response to vera and simon. *Cognitive Science*, *17*(1), 49–59. https://doi.org/10.1207/s15516709cog1701_3

Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic models of cognition: Exploring representations and inductive biases. *Trends In Cognitive Sciences*, *14*(8), 357–364. https://doi.org/10.1016/j.tics.2010.05.004

Griffiths, T. L., Kemp, C., & Tenenbaum, J. B. (2008). Bayesian models of cognition. In R. Sun (Ed.), *The cambridge handbook of computational psychology*. Cambridge University Press.

Guillot, A., Collet, C., Nguyen, V. A., Malouin, F., Richards, C., & Doyon, J. (2009). Brain activity during visual versus kinesthetic imagery: An fMRI study. *Human Brain Mapping*, *30*(7), 2157–2172.

Gunzelmann, G., & Anderson, J. R. (2003). Problem solving: Increased planning with practice. *Cognitive Systems Research*, *4*(1), 57–76. https://doi.org/10.1016/S1389-0417(02)00073-6

Handley, S. J., Capon, A., Copp, C., & Harper, C. (2002). Conditional reasoning and the tower of hanoi: The role of spatial and verbal working memory. *British Journal of Psychology*, *93*(4), 501–518. https://doi.org/10.1348/000712602761381376

Hart, R. A., & Moore, G. T. (1973). The development of spatial cognition: A review. In R. M. Downs & D. Stea (Eds.), *Image & environment: Cognitive mapping and spatial behavior* (pp. 246–288). New Brunswick, NJ, US, AldineTransaction.

Hatfield, G., & Epstein, W. (1985). The status of the minimum principle in the theoretical analysis of visual perception. *Psychological Bulletin*, *97*(2), 155–186. https://doi.org/10.1037/0033-2909.97.2.155

Hauk, O., Johnsrude, I., & Pulvermüller, F. (2004). Somatotopic representation of action words in human motor and premotor cortex. *Neuron*, *41*(2), 301–307. https://doi.org/10.1016/S0896-6273(03)00838-9

Havas, D. A., Glenberg, A. M., Gutowski, K. A., Lucarelli, M. J., & Davidson, R. J. (2010). Cosmetic use of botulinum toxin-a affects processing of emotional language. *Psychological Science*, *21*(7), 895–900.

Havas, D. A., & Matheson, J. (2013). The functional role of the periphery in emotional language comprehension. *Frontiers in Psychology*, *4*, 1–16.

Hiatt, L. M., Trafton, J. G., Harrison, A., & Schultz, A. C. (2004). A cognitive model for spatial perspective taking (M. C. Lovett, P. Munro, C. D. Schunn, & C. Lebiere, Eds.). In M. C. Lovett, P. Munro, C. D. Schunn, & C. Lebiere (Eds.), *Proceedings of the sixth international conference on cognitive modeling*.

Hingston, P. (2009). A turing test for computer game bots. *IEEE Transactions on Computational Intelligence and AI in Games*, *1*(3), 169–186. https://doi.org/10.1109/TCIAIG.2009.2032534

Hirtle, S. C., & Jonides, J. (1985). Evidence of hierarchies in cognitive maps. *Memory & Cognition*, *13*(3), 208–217.

Hölscher, C., Meilinger, T., Vrachliotis, G., Brösamle, M., & Knauff, M. (2006). Up the down staircase: Wayfinding strategies in multi-level buildings. *Journal of Environmental Psychology*, *26*(4), 284–299. https://doi.org/10.1016/j.jenvp.2006.09.002

Hyun, J.-S., & Luck, S. J. (2007). Visual working memory as the substrate for mental rotation. *Psychonomic bulletin & review*, *14*(1), 154–8.

Izard, V., & Dehaene, S. (2008). Calibrating the mental number line. *Cognition*, *106*(3), 1221–1247. https://doi.org/10.1016/j.cognition.2007.06.004

Jaderberg, M., Czarnecki, W. M., Dunning, I., Marris, L., Lever, G., Garcia Castaneda, A., Beattie, C., Rabinowitz, N. C., Morcos, A. S., Ruderman, A., Sonnerat, N., Green, T., Deason, L., Leibo, J. Z., Silver, D., Hassabis, D., Kavukcuoglu, K., & Graepel, T. (2018). Human-level performance in first-person multiplayer games with population-based deep reinforcement learning. *arXiv e-prints*, 1–42. https://ui.adsabs.harvard.edu/abs/2018arXiv180701281J

Jahn, G., Johnson-Laird, P. N., & Knauff, M. (2005). Reasoning about consistency with spatial mental models: Hidden and obvious indeterminacy in spatial descriptions. In C. Freksa, M. Knauff, B. Krieg-Brückner, B. Nebel, & T. Barkowsky (Eds.), *Spatial cognition iv. reasoning, action, interaction* (pp. 165–180). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-32255-9_10

Jahn, G., Knauff, M., & Johnson-Laird, P. N. (2007). Preferred mental models in reasoning about spatial relations. *Memory & Cognition*, *35*(8), 2075–2087. https://doi.org/10.3758/BF03192939

Johnson-Laird, P. N. (1972). The three-term series problem. *Cognition*, *1*(1), 57–82.

Johnson-Laird, P. N. (1989). Mental models. In M. Posner (Ed.), *Foundations of cognitive science*. Cambridge, MA: Bradford, MIT Press.

Jones, G., Ritter, F. E., & Wood, D. J. (2000). Using a cognitive architecture to examine what develops. *Psychological Science*, *11*(2), 93–100. https://doi.org/10.1111/1467-9280.00222

Kendall, G., Parkes, A. J., & Spoerer, K. (2008). A survey of np-complete puzzles. *ICGA Journal*, *31*(1), 13–34.

Kessler, K., & Thomson, L. A. (2010). The embodied nature of spatial perspective taking: Embodied transformation versus sensorimotor interference. *Cognition*, *114*(1), 72–88. https://doi.org/10.1016/j.cognition.2009.08.015

Kieras, D. E., & Meyer, D. E. (1997). An overview of the epic architecture for cognition and performance with application to human-computer interaction. *Human-Computer Interaction*, *12*(4), 391–438. https : / / doi . org / 10 . 1207 / s15327051hci1204_4

Kirsh, D. (1995). The intelligent use of space. *Artificial Intelligence*, *73*(1–2), 31–68.

Knauff, M., Rauh, R., & Schlieder, C. (1995). Preferred mental models in qualitative spatial reasoning: A cognitive assessment of allen's calculus, In *Proceedings of the seventeenth annual conference of the cognitive science society: July 22-25, 1995, university of pittsburgh*, Mahwah, NJ, Lawrence Erlbaum Associates. Psychology Press.

Kopp, B., Rist, F., & Mattler, U. (1996). N200 in the flanker task as a neurobehavioral tool for investigating executive control. *Psychophysiology*, *33*(3), 282–294. https://doi.org/10.1111/j.1469-8986.1996.tb00425.x

Kosslyn, S. M., Ganis, G., & Thompson, W. L. (2006). Mental imagery and the human brain. In Q. Jing, M. R. Rosenzweig, G. D'Ydewalle, H. Zhang, H.-C. Chen, & K. Zhang (Eds.), *Progress in psychological science around the world, vol 1: Neural, cognitive and developmental issues.* (pp. 195–209). New York, Psychology Press.

Kotovsky, K., Hayes, J. R., & Simon, H. A. (1985). Why are some problems hard? evidence from tower of hanoi. *Cognitive Psychology*, *17*(2), 248–294. https://doi.org/10.1016/0010-0285(85)90009-X

Kotovsky, K., & Simon, H. A. (1990). What makes some problems really hard: Explorations in the problem space of difficulty. *Cognitive Psychology*, *22*(2), 143–183. https://doi.org/10.1016/0010-0285(90)90014-U

Kuhn, G., Amlani, A. A., & Rensink, R. A. (2008). Towards a science of magic. *Trends in Cognitive Sciences*, *12*(9), 349–354. https://doi.org/10.1016/j.tics.2008.05.008

Kuipers, B. (1978). Modeling spatial knowledge. *Cognitive Science*, *2*(2), 129–153. https://doi.org/10.1207/s15516709cog0202_3

Kuipers, B. (1983). The cognitive map: Could it have been any other way?, In *Spatial orientation: Theory, research, and application*, Plenum Press.

Kuipers, B. (2000). The spatial semantic hierarchy. *Artificial Intelligence*, *119*(1-2), 191–233.

Kunze, L., & Beetz, M. (2017). Envisioning the qualitative effects of robot manipulation actions using simulation-based projections. *Artificial Intelligence*, *247*. https://doi.org/10.1016/j.artint.2014.12.004

Laird, J. E. (2008). Extending the soar cognitive architecture (P. Wang, B. Goertzel, & S. Franklin, Eds.). In P. Wang, B. Goertzel, & S. Franklin (Eds.), *Proceedings of the first conference on artificial general intelligence*, IOS Press.

Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, *11*(1), 65–100. https://doi.org/10.1111/j.1551-6708.1987.tb00863.x

Lee, D. N., & Reddish, P. E. (1981). Plummeting gannets: A paradigm of ecological optics. *Nature*, *293*(5830), 293–294. https://doi.org/10.1038/293293a0
10.1038/293293a0

Leeuwen, M. V. (2005). Questions for the dynamicist: The use of dynamical systems theory in the philosophy of cognition. *Minds and Machines*, *15*(3), 271–333. https://doi.org/10.1007/s11023-004-8339-2

Levine, M., Marchon, I., & Hanley, G. (1984). The placement and misplacement of you-are-here maps. *Environment and Behavior*, *16*(2), 139–157. https://doi.org/10.1177/0013916584162001

Levinson, S. C. (1996). Frames of reference and molyneuxs question: Crosslinguistic evidence. In P. Bloom, M. Peterson, L. Nadel, & M. Garrett (Eds.), *Language and space* (pp. 109–169). Cambridge, MA, MIT Press.

Levinson, S. C. (1998). Studying spatial conceptualization across cultures: Anthropology and cognitive science. *Ethos*, *26*(1), 7–24. https://doi.org/10.1525/eth.1998.26.1.7

Lewis, C. (1982). *Using the "thinking-aloud" method in cognitive interface design* (Research Report RC 9265). IBM T.J. Watson Research Center. Yorktown Heights, NY.

Liberman, A. M., Cooper, F. S., Shankweiler, D. P., & Studdert-Kennedy, M. (1967). Perception of the speech code. *Psychological Review*, *74*(6), 431–461. https://doi.org/10.1037/h0020279

Liesefeld, H. R., & Zimmer, H. D. (2013). Think spatial: The representation in mental rotation is nonvisual. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(1), 167–182. https://doi.org/10.1037/a0028904

Lovett, A., & Schultheis, H. (2014). Modeling spatial abstraction during mental rotation (P. Bello, M. McShane, M. Guarini, & B. Scassellati, Eds.). In P. Bello, M. McShane, M. Guarini, & B. Scassellati (Eds.), *Proceedings of the 36th annual meeting of the cognitive science society*.

MacGregor, J. N., & Ormerod, T. C. (1996). Human performance on the traveling salesman problem. *Perception & Psychophysics*, *58*(4), 527–539.

Madl, T., Chen, K., Montaldi, D., & Trappl, R. (2015). Computational cognitive models of spatial memory in navigation space: A review. *Neural Networks*, *65*, 18–43. https://doi.org/10.1016/j.neunet.2015.01.002

Majid, A., Bowerman, M., Kita, S., Haun, D. B. M., & Levinson, S. C. (2004). Can language restructure cognition? the case for space. *Trends in Cognitive Sciences*, *8*(3), 108–114. https://doi.org/10.1016/j.tics.2004.01.003

Marewski, J. N., & Mehlhorn, K. (2011). Using the ACT-R architecture to specify 39 quantitative process models of decision making. *Judgment and Decision Making*, *6*(6), 439–519.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco, CA, W.H. Freeman.

Marr, D., & Nishihara, H. K. (1978). Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, *200*(1140), 269–294.

Martin, K. A., Moritz, S. E., & Hall, C. R. (1999). Imagery use in sport: A literature review and applied model. *The sport psychologist*, *13*(3), 245–268.

McBeath, M. K., Shaffer, D. M., & Kaiser, M. K. (1995). How baseball outfielders determine where to run to catch fly balls. *Science*, *268*, 569–573.

McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T. T., Seidenberg, M. S., & Smith, L. B. (2010). Letting structure emerge: Connectionist and dynamical systems approaches to cognition. *Trends in Cognitive Sciences*, *14*(8), 348–356. https://doi.org/10.1016/j.tics.2010.06.002

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. an account of basic findings. *Psychological Review*, *88*(5), 375–407. https://doi.org/10.1037/0033-295x.88.5.375

McGurk, H., & MacDonald, J. (1976). Hearing lips and seeing voices. *Nature*, *264*(5588), 746–748. https://doi.org/10.1038/264746a0

Miller, G. A. (2003). The cognitive revolution: A historical perspective. *Trends in Cognitive Sciences*, *7*(3), 141–144. https://doi.org/10.1016/S1364-6613(03)00029-9

Moeller, K., Fischer, U., Link, T., Wasner, M., Huber, S., Cress, U., & Nuerk, H.-C. (2012). Learning and development of embodied numerosity. *Cognitive Processing*, *13*(Supplement 1), 271–274.

Mordvintsev, A., Olah, C., & Tyka, M. (2015). *Inceptionism: Going deeper into neural networks*. https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html
essay from his webpage

Moreno, R., & Mayer, R. E. (1999). Cognitive principles of multimedia learning: The role of modality and contiguity. *Journal of Educational Psychology*, *91*(2), 358–368. https://doi.org/10.1037/0022-0663.91.2.358

Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology*, *9*(3), 353–383. https://doi.org/10.1016/0010-0285(77)90012-3

Necker, L. A. (1832). LXI. observations on some remarkable optical phænomena seen in switzerland; and on an optical phænomenon which occurs on viewing a figure of a crystal or geometrical solid. *Philosophical Magazine Series 3*, *1*(5), 329–337. https://doi.org/10.1080/14786443208647909

Newell, A. (1980). Physical symbol systems. *Cognitive Science*, *4*(2), 135–183.

Newell, A., & Simon, H. A. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, *19*(3), 113–126. https://doi.org/10.1145/360018.360022

Noordzij, M. L., & Postma, A. (2005). Categorical and metric distance information in mental representations derived from route and survey descriptions. *Psychological research*, *69*(3), 221–232. https://doi.org/10.1007/s00426-004-0172-y

Noordzij, M. L., Zuidhoek, S., & Postma, A. (2006). The influence of visual experience on the ability to form spatial mental models based on route and survey descriptions. *Cognition*, *100*(2), 321–342. https://doi.org/10.1016/j.cognition.2005.05.006

Nyhus, E., & Barceló, F. (2009). The wisconsin card sorting test and the cognitive assessment of prefrontal executive functions: A critical update. *Brain and Cognition*, *71*(3), 437–451. https://doi.org/10.1016/j.bandc.2009.03.005

Olkun, S. (2003). Comparing computer versus concrete manipulatives in learning 2d geometry. *Journal of Computers in Mathematics and Science Teaching*, *22*(1), 43–56.

Palmer, S. E., Brooks, J. L., & Nelson, R. (2003). When does grouping happen? [Visual Gestalt Formation]. *Acta Psychologica*, *114*(3), 311–330. https://doi.org/10.1016/j.actpsy.2003.06.003

Pande, P., & Chandrasekharan, S. (2014). Eye-tracking in STEM education research: Limitations, experiences and possible extensions (Kinshuk & S. Murthy, Eds.). In Kinshuk & S. Murthy (Eds.), *Technology for Education (T4E), 2014 IEEE Sixth International Conference on*. https://doi.org/10.1109/T4E.2014.29

Pasteur, G. (1982). A classificatory review of mimicry systems. *Annual Review of Ecology and Systematics*, *13*(1), 169–199. https://doi.org/10.1146/annurev.es.13.110182.001125

Piccinini, G. (2000). Turing's rules for the imitation game. *Minds and Machines*, *10*(4), 573–582. https://doi.org/10.1023/A:1011246220923

Pinna, B., Brelstaff, G., & Spillmann, L. (2001). Surface color from boundaries: A new 'watercolor' illusion. *Vision Research*, *41*(20), 2669–2676. https://doi.org/10.1016/S0042-6989(01)00105-5

Purves, D., Beau Lotto, R., Mark Williams, S., Nundy, S., & Yang, Z. (2001). Why we see things the way we do: Evidence for a wholly empirical strategy of vision. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *356*(1407), 285–297. https://doi.org/10.1098/rstb.2000.0772

Pylyshyn, Z. (2002). Mental imagery: In search of a theory. *Behavioral and Brain Sciences*, *25*(02), http://journals.cambridge.org/article_S0140525X02000043, 157–182. https://doi.org/10.1017/S0140525X02000043

Pylyshyn, Z. (2003). Return of the mental image: Are there really pictures in the brain? *Trends in Cognitive Sciences*, *7*(3), 113–118. https://doi.org/10.1016/S1364-6613(03)00003-2

Ragni, M., Fangmeier, T., & Brüssow, S. (2010). Deductive spatial reasoning: From neurological evidence to a cognitive model (D. D. Salvucci & G. Gunzelmann, Eds.). In D. D. Salvucci & G. Gunzelmann (Eds.), *Proceedings of the 10th international conference on cognitive modeling*, Philadelphia, PA: Drexel University.

Randell, D. A., Cui, Z., & Cohn, A. G. (1992). A spatial logic based on regions and connection (B. Nebel, C. Rich, & W. Swartout, Eds.). In B. Nebel, C. Rich, & W. Swartout (Eds.), *Principles of knowledge representation and reasoning. procedings of the third international conference.* Los Altos, CA, Kaufmann, Morgan.

Rauh, R., Hagen, C., Knauff, M., Kuß, T., Schlieder, C., & Strube, G. (2005). Preferred and alternative mental models in spatial reasoning. *Spatial Cognition & Computation*, *5*(2), 239–269.

Raven, J. (2000). The raven's progressive matrices: Change and stability over culture and time. *Cognitive Psychology*, *41*(1), 1–48. https://doi.org/10.1006/cogp.1999.0735

Rayner, K., White, S. J., Johnson, R. L., & Liversedge, S. P. (2006). Raeding wrods with jubmled lettres: There is a cost. *Psychological Science*, *17*(3), 192–193. https://doi.org/10.1111/j.1467-9280.2006.01684.x

Reitman, J. S. (1976). Skilled perception in go: Deducing memory structures from inter-response times. *Cognitive Psychology*, *8*(3), 336–356. https://doi.org/10.1016/0010-0285(76)90011-6

Reitter, D., & Lebiere, C. (2010). A cognitive model of spatial path-planning. *Computational and Mathematical Organization Theory*, *16*(3), 220–245. https://doi.org/10.1007/s10588-010-9073-3

Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, *8*(5), 368–373. https://doi.org/10.1111/j.1467-9280.1997.tb00427.x

Rizzolatti, G., & Craighero, L. (2004). The mirror-neuron system. *Annual review of neuroscience*, *27*, 169–192. https://doi.org/10.1146/annurev.neuro.27.070203.144230

Robinson, H. (2016). Embodied cognition. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Winter 2016 ed.). https://plato.stanford.edu/archives/win2016/entries/dualism

Rock, I., & Palmer, S. (1990). The legacy of gestalt psychology. *Scientific American*, *263*(6), 84–90

Argue that grouping by Gestalt happens after early vusiual processing, thus grouping is based on the perceived structure of the real world rather than on the retina. p.87.

Ronsse, R., Puttemans, V., Coxon, J. P., Goble, D. J., Wagemans, J., Wenderoth, N., & Swinnen, S. P. (2010). Motor learning with augmented feedback: Modality-dependent behavioral and neural consequences. *Cerebral Cortex*, *21*(6), 1283–1294. https://doi.org/10.1093/cercor/bhq209

Rossignac, J. R., & Borrel, P. (1992). *Multi-resolution 3D approximations for rendering complex scenes* (Research Report). IBM Research Division; T.J. Watson Research Center. Yorktown Heights, NY.

Rumelhart, D. E., & Abrahamson, A. A. (1973). A model for analogical reasoning. *Cognitive Psychology*, *5*(1), 1–28. https://doi.org/10.1016/0010-0285(73)90023-6

Rusinkiewicz, S., & Levoy, M. (2000). Qsplat: A multiresolution point rendering system for large meshes, In *Proceedings of the 27th annual conference on Computer Graphics and Interactive Techniques*, New York, NY, ACM Press/Addison-Wesley Publishing Co. https://doi.org/10.1145/344779.344940

Salomon, G., Perkins, D. N., & Globerson, T. (1991). Partners in cognition: Extending human intelligence with intelligent technologies. *Educational Researcher*, *20*(3), 2–9.

Sandin, L. (2020). *The SSDF rating list*. http://ssdf.bosjo.net/long.txt

Scaife, M., & Rogers, Y. (1996). External cognition: How do graphical representations work? *International Journal of Human-Computer Studies*, *45*(2), 185–213. https://doi.org/10.1006/ijhc.1996.0048

Schaeffer, J., Burch, N., Björnsson, Y., Kishimoto, A., Müller, M., Lake, R., Lu, P., & Sutphen, S. (2007). Checkers is solved. *Science*, *317*(5844), 1518–1522.

Schaeffer, J., Lake, R., Lu, P., & Bryant, M. (1995). CHINOOK - the world man-machine checkers champion. *AI Magazine*, *17*(1), 21–29.

Schöner, G. (2008). Dynamical systems approaches to cognition. In R. Sun (Ed.), *The cambridge handbook of computational psychology*. Cambridge University Press.

Schultheis, H., & Barkowsky, T. (2011). Casimir: An architecture for mental spatial knowledge processing. *Topics in Cognitive Science*, *3*(4), 778–795. https://doi.org/10.1111/j.1756-8765.2011.01151.x

Schultheis, H., Bertel, S., & Barkowsky, T. (2014). Modeling mental spatial reasoning about cardinal directions. *Cognitive Science*, *38*(8), 1521–1561. https://doi.org/10.1111/cogs.12125

Schultheis, H., & Wienemann, R. (2019, September 2–4). Is space special? the role of spatial information in analogical reasoning, problem solving, and decision making [poster session]. EuroCogSci 2019 Situated minds and flexible cognition, Ruhr-Universität Bochum, Germany.

Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, *3*(3), 417–424. https://doi.org/10.1017/S0140525X00005756

Sebbane, M. (2001). Board games from canaan in the early and intermediate bronze ages and the origin of the egyptian senet game. *Tel Aviv*, *28*(2), 213–230. https://doi.org/10.1179/tav.2001.2001.2.213

Selfridge, O. G. (1955). Pattern recognition and modern computers, In *Proceedings of the march 1-3, 1955, western joint computer conference*, Los Angeles, California, ACM. https://doi.org/10.1145/1455292.1455310

Shah, H., & Warwick, K. (2010). Testing turing's five minutes, parallel-paired imitation game. *Kybernetes*, *39*(3), 449–465. https://doi.org/10.1108/03684921011036178

Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *298*(1089), 199–209.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, *27*(3), 379–423.

Shepard, R. N. (1984). Ecological constraints on internal representation: Resonant kinematics of perceiving, imagining, thinking, and dreaming. *Psychological Review*, *91*(4), 417–447. https://doi.org/10.1037/0033-295x.91.4.417

Shepard, R. N. (1994). Perceptual-cognitive universals as reflections of the world. *Psychonomic Bulletin & Review*, *1*(1), 2–28. https://doi.org/10.3758/BF03200759

Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, *171*(3972), 701–703. https://doi.org/10.1126/science.171.3972.701

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016). Mastering the game of go with deep neural networks and tree search. *Nature*, *529*(7587), 484–489. https://doi.org/10.1038/nature16961

Sima, J. F., Lindner, M., Schultheis, H., & Barkowsky, T. (2010). Eye movements reflect reasoning with mental images but not with mental models in orientation knowledge tasks (C. Hölscher, T. Shipley, M. Olivetti Belardinelli, J. Bateman, & N. Newcombe, Eds.). In C. Hölscher, T. Shipley, M. Olivetti Belardinelli, J. Bateman, & N. Newcombe (Eds.), *Spatial cognition VII*, Berlin, Heidelberg, Germany, Springer-Verlag.

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, *69*(1), 99–118.

Simon, H. A. (1972). Theories of bounded rationality. In C. B. McGuire & R. Radner (Eds.), *Decision and organization* (pp. 161–176). North-Holland.

Simon, H. A. (1978). On the forms of mental representation. In C. Savage (Ed.), *Minnesota studies in the philosophy of science. Vol. IX: Perception and cognition: Issues in the foundations of psychology* (pp. 3–18). Minneapolis, University of Minnesota Press.

Simon, H. A., & Newell, A. (1958). Heuristic problem solving: The next advance in operations research. *Operations Research*, *6*(1), 1–10.

Simon, J. R. (1990). The effects of an irrelevant directional cue on human information processing. In R. W. Proctor & T. G. Reeve (Eds.), *Advances in psychology* (pp. 31–86). Elsevier. https://doi.org/10.1016/S0166-4115(08)61218-2

Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception*, *28*(9), 1059–1074. https://doi.org/10.1068/p281059

Simons, D. J., & Rensink, R. A. (2005). Change blindness: Past, present, and future. *Trends in Cognitive Sciences*, *9*(1), 16–20. https://doi.org/10.1016/j.tics.2004.11.006

Sloman, A. (1971). Interactions between philosophy and artificial intelligence: The role of intuition and non-logical reasoning in intelligence. *Artificial Intelligence*, *2*(3), 209–225. https://doi.org/10.1016/0004-3702(71)90011-7

Sloman, A. (1975). Afterthoughts on analogical representations, In *Proceedings of the 1975 workshop on theoretical issues in natural language processing*, Cambridge, Massachusetts, Association for Computational Linguistics. https://doi.org/10.3115/980190.980235

Smith, E. R., & Semin, G. R. (2007). Situated social cognition. *Current Directions in Psychological Science*, *16*(3), 132–135. https://doi.org/10.1111/j.1467-8721.2007.00490.x

Steffenhagen, F., Albrecht, R., & Ragni, M. (2014). Automatic identification of human strategies by cognitive agents. In C. Lutz & M. Thielscher (Eds.), *Lecture Notes in Computer Science: Vol. 8736: KI 2014: Advances in Artificial Intelligence* (pp. 62–67). Springer International Publishing. https://doi.org/10.1007/978-3-319-11206-0_7

Stenning, K., & Lemon, O. (2001). Aligning logical and psychological perspectives on diagrammatic reasoning. *Artificial Intelligence Review*, *15*(1), 29–62. https://doi.org/10.1023/A:1006617525134

Stenning, K., & Tobin, R. (1997). Assigning information to modalities: Comparing graphical treatments of the syllogism. In E. Ejerhed & S. Lindström (Eds.), *Logic, action and cognition: Essays in philosophical logic* (pp. 211–228). Dordrecht, Springer Netherlands. https://doi.org/10.1007/978-94-011-5524-3_10

Stevens, A., & Coupe, P. (1978). Distortions in judged spatial relations. *Cognitive Psychology*, *10*(4), 422–437. https://doi.org/10.1016/0010-0285(78)90006-3

Stoianov, I., Kramer, P., Umiltà, C., & Zorzi, M. (2008). Visuospatial priming of the mental number line. *Cognition*, *106*(2), 770–779. https://doi.org/10.1016/j.cognition.2007.04.013

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, *18*(6), 643–662. https://doi.org/10.1037/h0054651

Sun, R. (2008). Introduction to computational cognitive modeling. In *The cambridge handbook of computational psychology*. Cambridge University Press.

Szabó, B. T., Denham, S. L., & Winkler, I. (2016). Computational models of auditory scene analysis: A review. *Frontiers in Neuroscience*, *10*, 1–16. https://doi.org/10.3389/fnins.2016.00524

Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences*, *24*(1), 1–34. https://doi.org/10.1017/S0140525X01003910

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological review*, *55*(4), 189–208.

Tooby, J., & Cosmides, L. (2005). Evolutionary psychology: Conceptual foundations. In D. M. Buss (Ed.), *The handbook of evolutionary psychology* (pp. 5–67). Hoboken, NJ, Wiley.

Tresilian, J. R. (1999). Visually timed action: Time-out for ‘tau’? *Trends in Cognitive Sciences*, *3*(8), 301–310.

Turing, A. M. (1937). On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, *s2-42*(1), 230–265. https://doi.org/10.1112/plms/s2-42.1.230

Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, *59*(236), 433–460. https://doi.org/10.1093/mind/LIX.236.433

Turner, A. (2009). The role of angularity in route choice: An analysis of motorcycle courier gps traces, In *Proceedings of the 9th international conference on spatial information theory*, Springer Berlin Heidelberg New York.

Tversky, B. (1981). Distortions in memory for maps. *Cognitive Psychology*, *13*(3), 407–433. https://doi.org/10.1016/0010-0285(81)90016-5

Tversky, B. (2009). Spatial cognition: Embodied and situated. In P. Robbins & M. Aydede (Eds.), *The cambridge handbook of situated cognition* (pp. 201–217). NY, Cambridge University Press.

Tversky, B., Bauer Morrison, J., Franklin, N., & Bryant, D. J. (1999). Three spaces of spatial cognition. *The Professional Geographer*, *51*(4), 516–524. https://doi.org/10.1111/0033-0124.00189

Tversky, B., & Hard, B. M. (2009). Embodied and disembodied cognition: Spatial perspective-taking. *Cognition*, *110*(1), 124–129. https://doi.org/10.1016/j.cognition.2008.10.008

Tversky, B., Heiser, J., Lee, P., & Daniel, M.-P. (2009). Explanations in gesture, diagram, and word. *Spatial language and dialogue*, *1*(9), 119–132.

van den Herik, H., Uiterwijk, J. W., & van Rijswijck, J. (2002). Games solved: Now and in the future. *Artificial Intelligence*, *134*(1–2), 277–311
state space of games among other things.

van Geert, P. (1998). A dynamic systems model of basic developmental mechanisms: Piaget, vygotsky, and beyond. *Psychological Review*, *105*(4), 634–677. https://doi.org/10.1037/0033-295x.105.4.634-677

van Gelder, T. (1995). What might cognition be, if not computation? *The Journal of Philosophy*, *92*(7), 345–381.

van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, *21*(5), 615–628.

van Gelder, T., & Port, R. F. (1995). It's about time: An overview of the dynamical approach to cognition. In R. F. Port & T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition* (pp. 1–43). The MIT Press.

van Rijn, H., van Someren, M., & van der Maas, H. (2003). Modeling developmental transitions on the balance scale task. *Cognitive Science*, *27*(2), 227–257. https://doi.org/10.1016/S0364-0213(02)00116-7

Veksler, V. D., Gray, W. D., & Schoelles, M. J. (2013). Goal-proximity decision-making. *Cognitive Science*, *37*(4), 757–774. https://doi.org/10.1111/cogs.12034

Velhagen, K., & Broschmann, D. (1995). *Tafeln zur prüfung des farbensinnes* (30th). Thieme.

Vernon, D., Metta, G., & Sandini, G. (2007). A survey of artificial cognitive systems: Implications for the autonomous development of mental capabilities in computational agents. *IEEE Transactions on Evolutionary Computation*, *11*(2), 151–180. https://doi.org/10.1109/TEVC.2006.890274

Vinyals, O., Babuschkin, I., Chung, J., Mathieu, M., Jaderberg, M., Czarnecki, W. M., Dudzik, A., Huang, A., Georgiev, P., Powell, R., Ewalds, T., Horgan, D., Kroiss, M., Danihelka, I., Agapiou, J., Oh, J., Dalibard, V., Choi, D., Sifre, L., … Silver, D. (2019). *AlphaStar: Mastering the Real-Time Strategy Game StarCraft II*. https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/

von Ahn, L., Blum, M., Hopper, N. J., & Langford, J. (2003). CAPTCHA: Using hard ai problems for security. In E. Biham (Ed.), *Advances in cryptology — eurocrypt 2003: International conference on the theory and applications of cryptographic techniques, warsaw, poland, may 4–8, 2003 proceedings* (pp. 294–311). Berlin, Heidelberg, Springer Berlin Heidelberg. https://doi.org/10.1007/3-540-39200-9_18

Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., & von der Heydt, R. (2012). A century of gestalt psychology in visual perception: I. perceptual grouping and figure–ground organization. *Psychological Bulletin*, *138*(6), 1172–1217. https://doi.org/10.1037/a0029333

Wagemans, J., Feldman, J., Gepshtein, S., Kimchi, R., Pomerantz, J. R., van der Helm, P. A., & van Leeuwen, C. (2012). A century of gestalt psychology in visual perception: II. conceptual and theoretical foundations. *Psychological Bulletin*, *138*(6), 1218–1252. https://doi.org/10.1037/a0029334

Wai, J., Lubinski, D., & Benbow, C. P. (2009). Spatial ability for STEM domains: Aligning over 50 years of cumulative psychological knowledge solidifies its importance.aligning over 50 years of cumulative psychological knowledge solidifies its importance. *Journal of Educational Psychology*, *101*(4), 817–835. https://doi.org/10.1037/a0016127

Warren, W. H. (1984). Perceiving affordances: Visual guidance of stair climbing. *Journal of Experimental Psychology: Human Perception and Performance*, *10*(5), 683–703. https://doi.org/10.1037/0096-1523.10.5.683

Wesp, R., Hesse, J., Keutmann, D., & Wheaton, K. (2001). Gestures maintain spatial imagery. *The American journal of psychology*, *114*(4), 591–600.

Wilson, F. A. W., Scalaidhe, S. P. Ó., & Goldman-Rakic, P. S. (1993). Dissociation of object and spatial processing domains in primate prefrontal cortex. *Science*, *260*(5116), 1955–1958.

Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, *9*(4), 625–636.

Wilson, R. A., & Clark, A. (2009). How to situate cognition: Letting nature take its course. In P. Robbins & M. Aydede (Eds.), *The cambridge handbook of situated cognition* (pp. 55–77). Cambridge University Press.

Wilson, R. A., & Foglia, L. (2015). Embodied cognition. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Spring 2016 ed.). http://plato.stanford.edu/archives/spr2016/entries/embodied-cognition/

Wolpert, D. (2011). The real reason for brains [Video file]. https://www.ted.com/talks/daniel_wolpert_the_real_reason_for_brains

Wundt, W. M. (1874). *Grundzüge der physiologischen psychology*. Leipzig, Germany, Engelmann.

Yang, J., & Schwaninger, A. (2011). Turn right or turn left? heuristic of adhering to the direction of destination. *Applied Cognitive Psychology*, *25*(5), 703–707. https://doi.org/10.1002/acp.1741

Zhang, J. (1997). The nature of external representations in problem solving. *Cognitive Science*, *21*(2), 179–217. https://doi.org/10.1016/S0364-0213(99)80022-6

Appendices

Appendix A.

Tic-tac-toe experiment

Stimuli of the tic-tac-toe experiment as well as supplementary material is available for use at: https://github.com/wienemann/experiments

A.1. Experiemental protocol

1. Welcome
2. Ask to sign an informed consent form
3. Give the participant the instructions
4. Conduct color vision test
5. If they have no further questions let them start the experimental program
6. When done give them the demographics questionnaire
7. Give the participant the compensation and let them sign for it.

A.2. Post-test questionnaire

cognitivesystems

Experiment zum Spielverhalten

Abschließender Fragebogen

Geschlecht: O Weiblich
O Männlich

Alter: _____

Nationalität: O Deutsch
O ______________

Wie hast du von dem Experiment gehört?
O [psycho]-Verteiler
O Stud.IP
O Freunde
O ___________________

Studienfach: ____________________

Was für eine Strategie hast du genutzt?

Kennst du das Spiel Tic-Tac-Toe (auch bekannt als Drei Gewinnt, Käsekästchen oder XXO)? Wenn ja (wie) hat es dir bei diesem Experiment geholfen?
O nein O ja:

Sonstige Anmerkungen zum Experiment:

Figure A.1.: Post-test questionnaire

A.3. Novel Stimuli

A.3.1. Condition 2; randomized placement

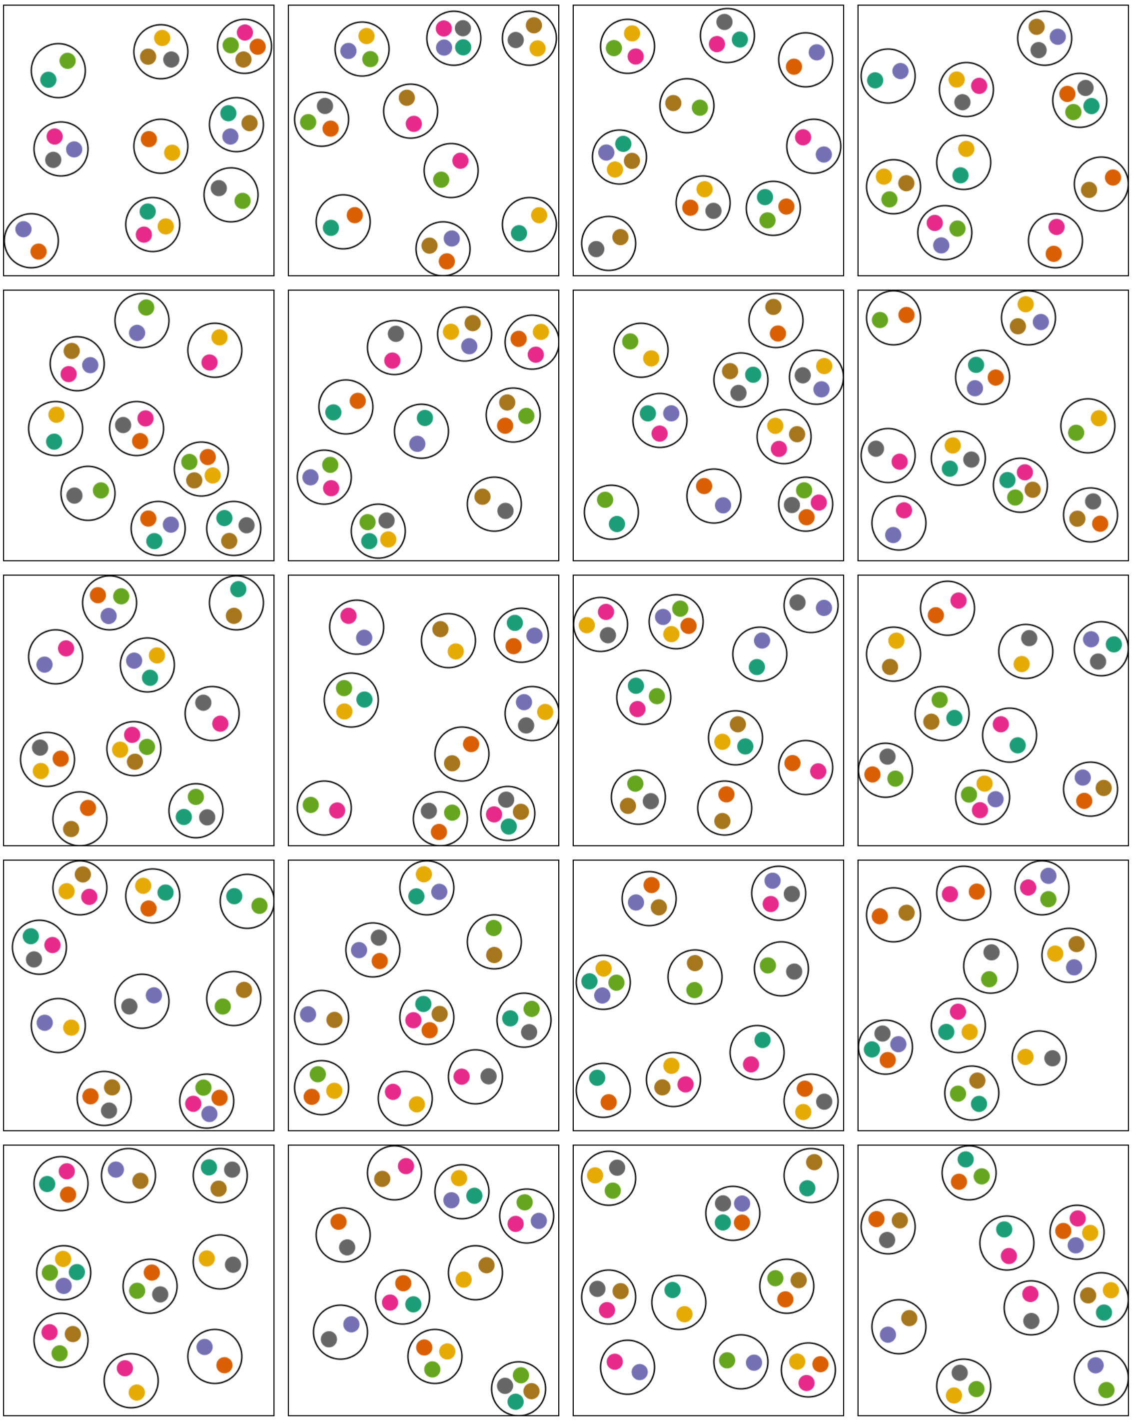

A.3.2. Condition 3; interference placement

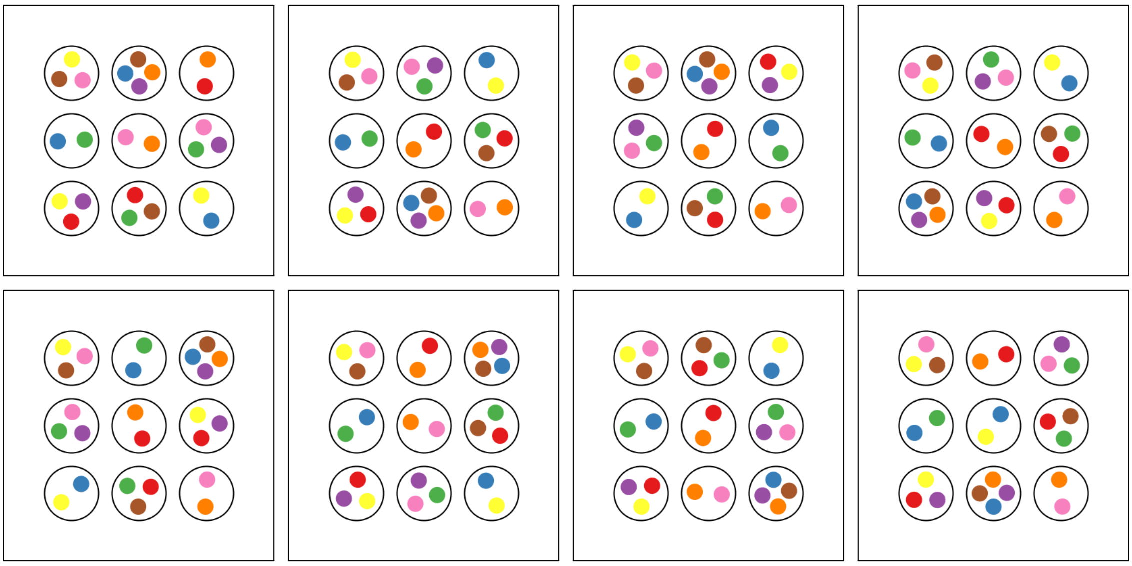

Appendix B.

Card sorting experiment all stimuli

Stimuli of the card sorting experiment among others are available for use at:
https://github.com/wienemann/experiments

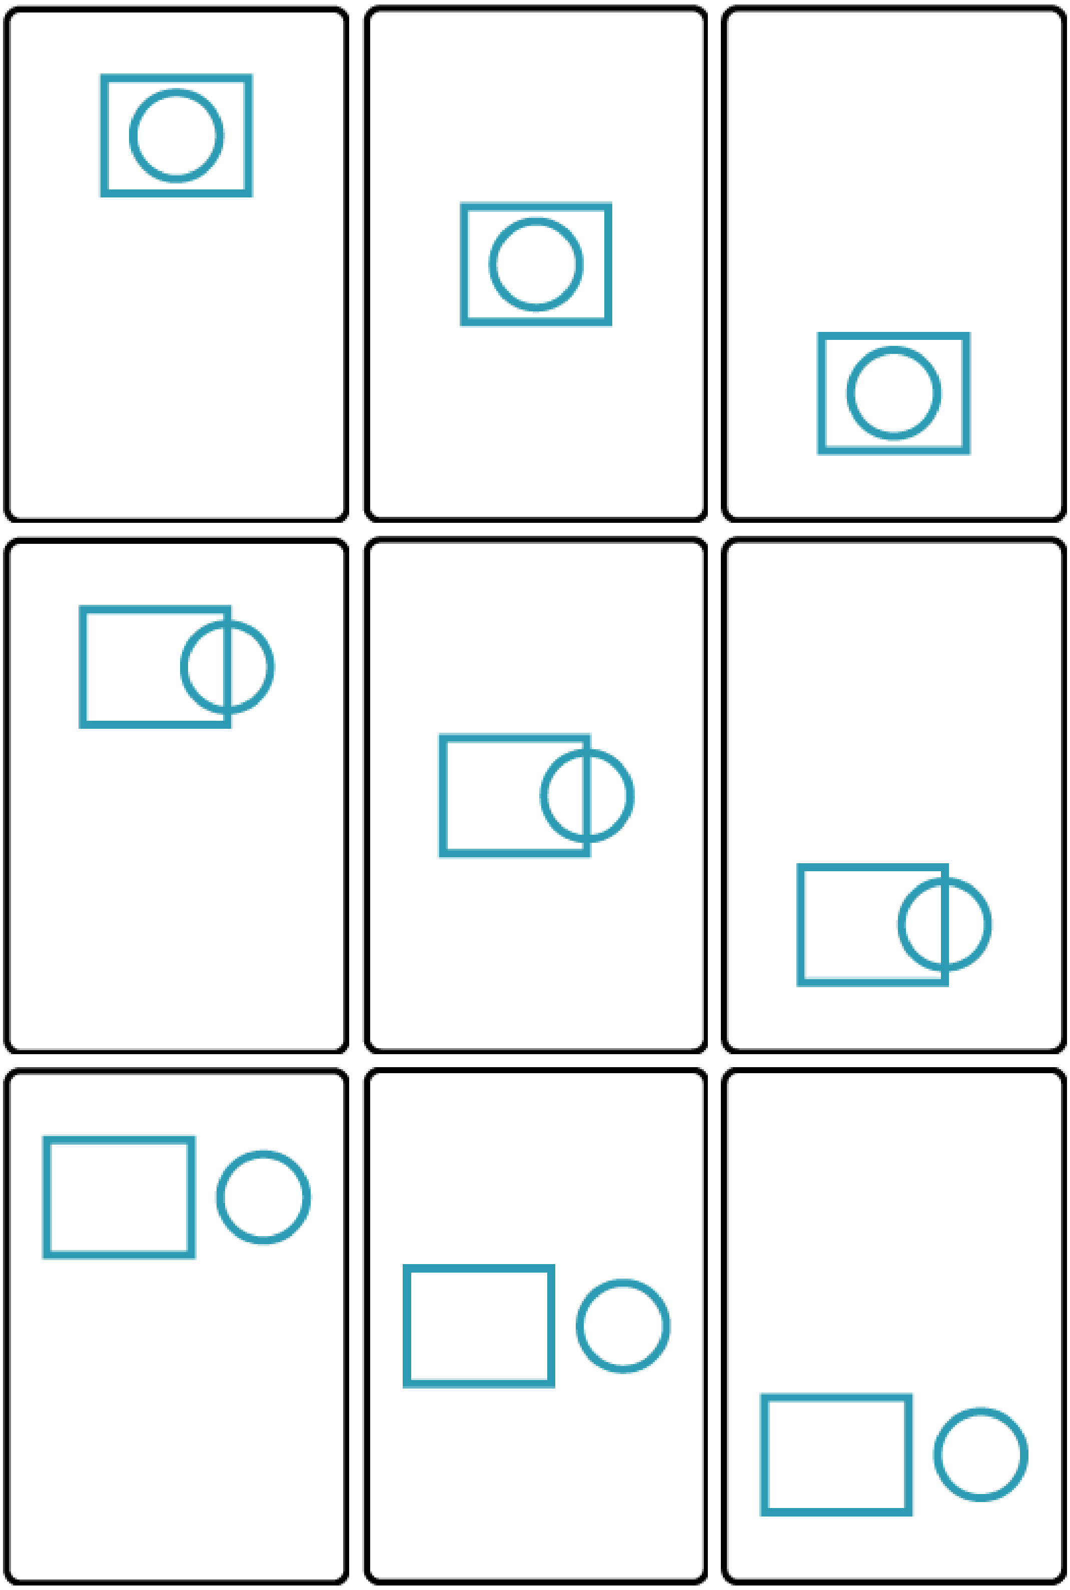

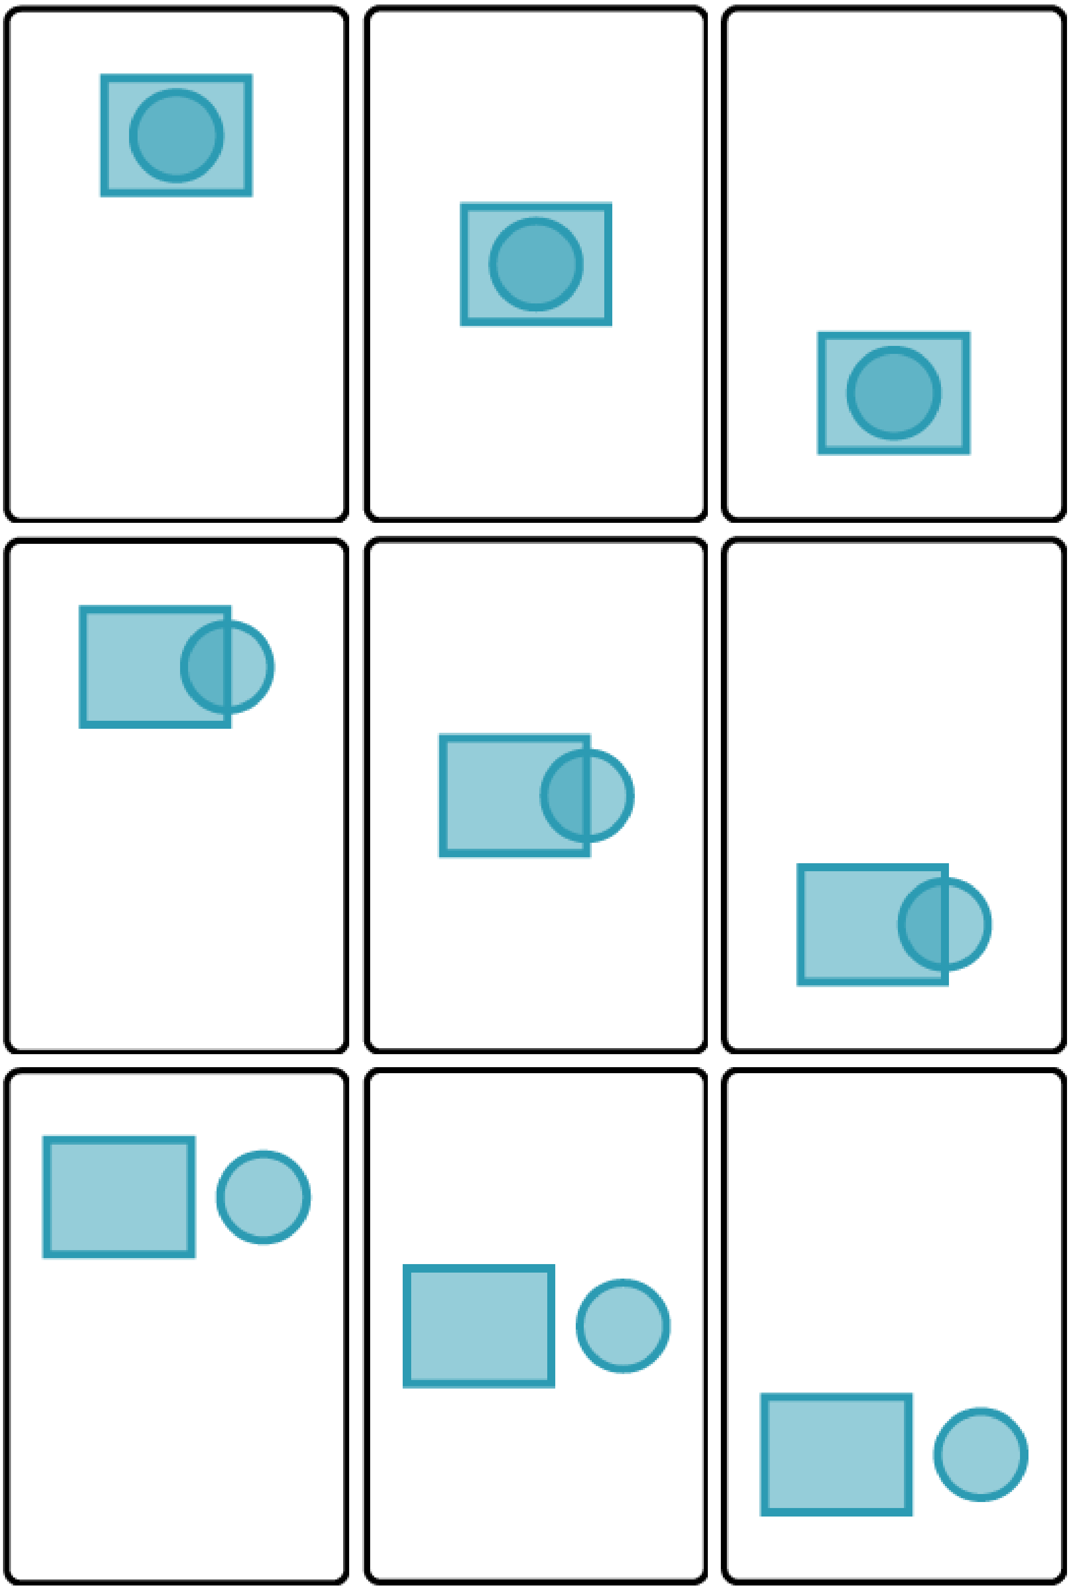

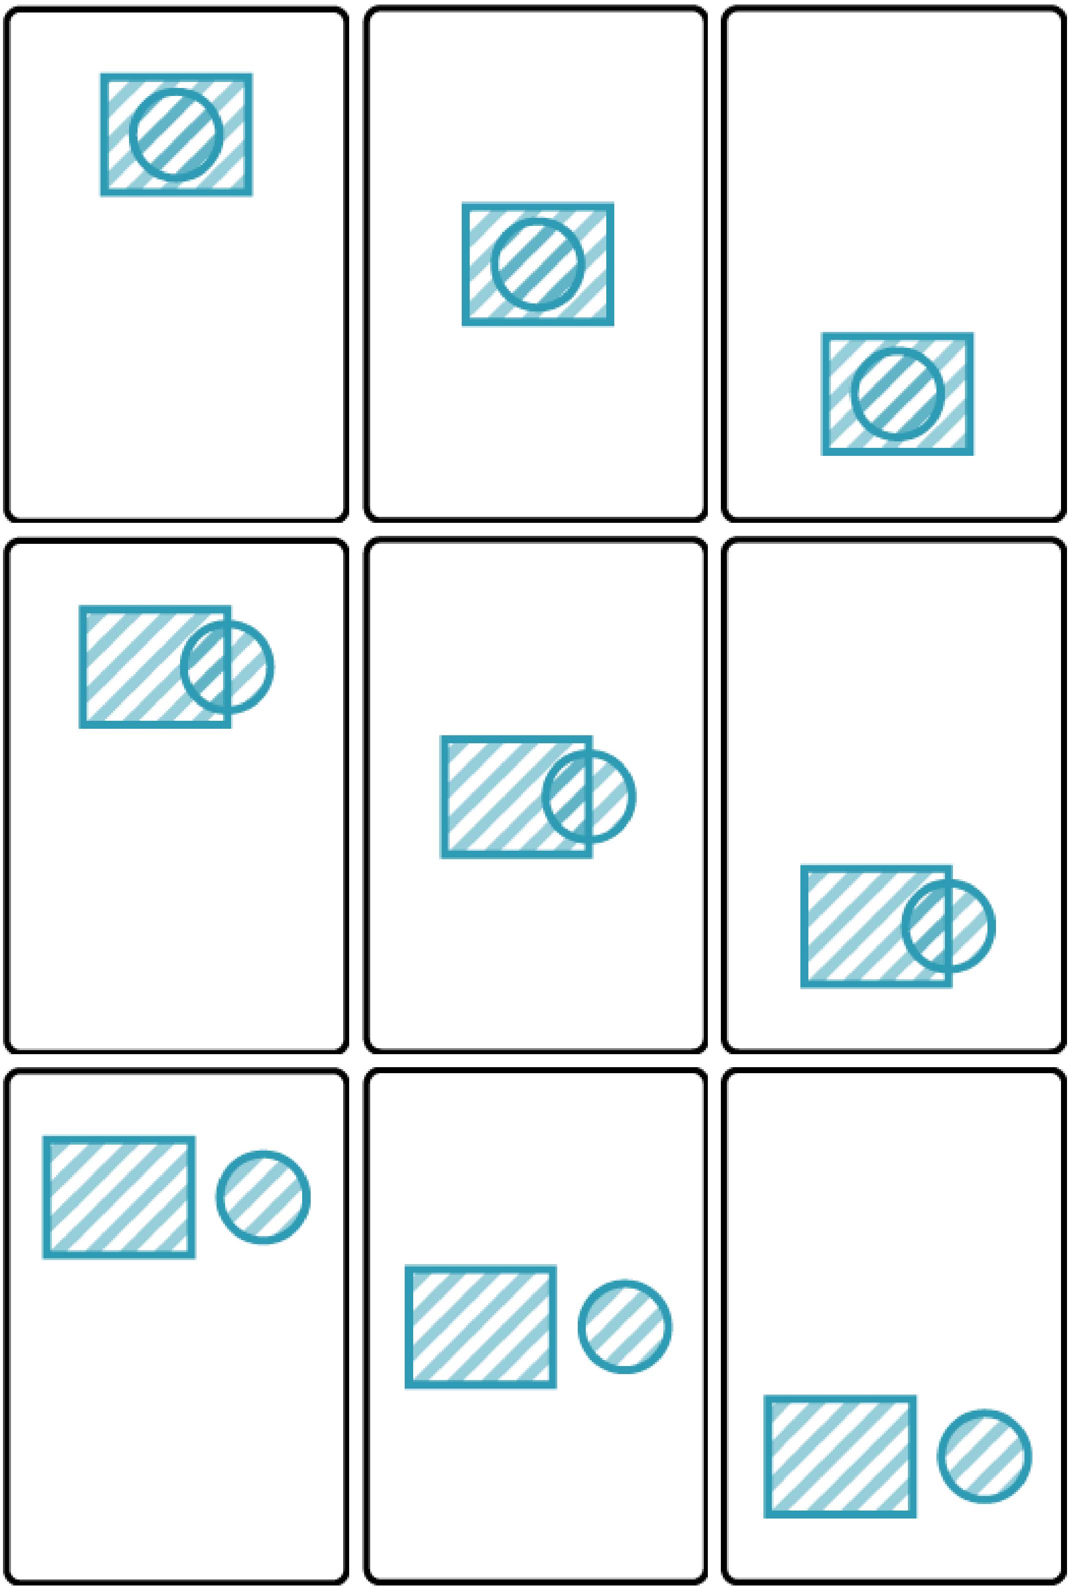

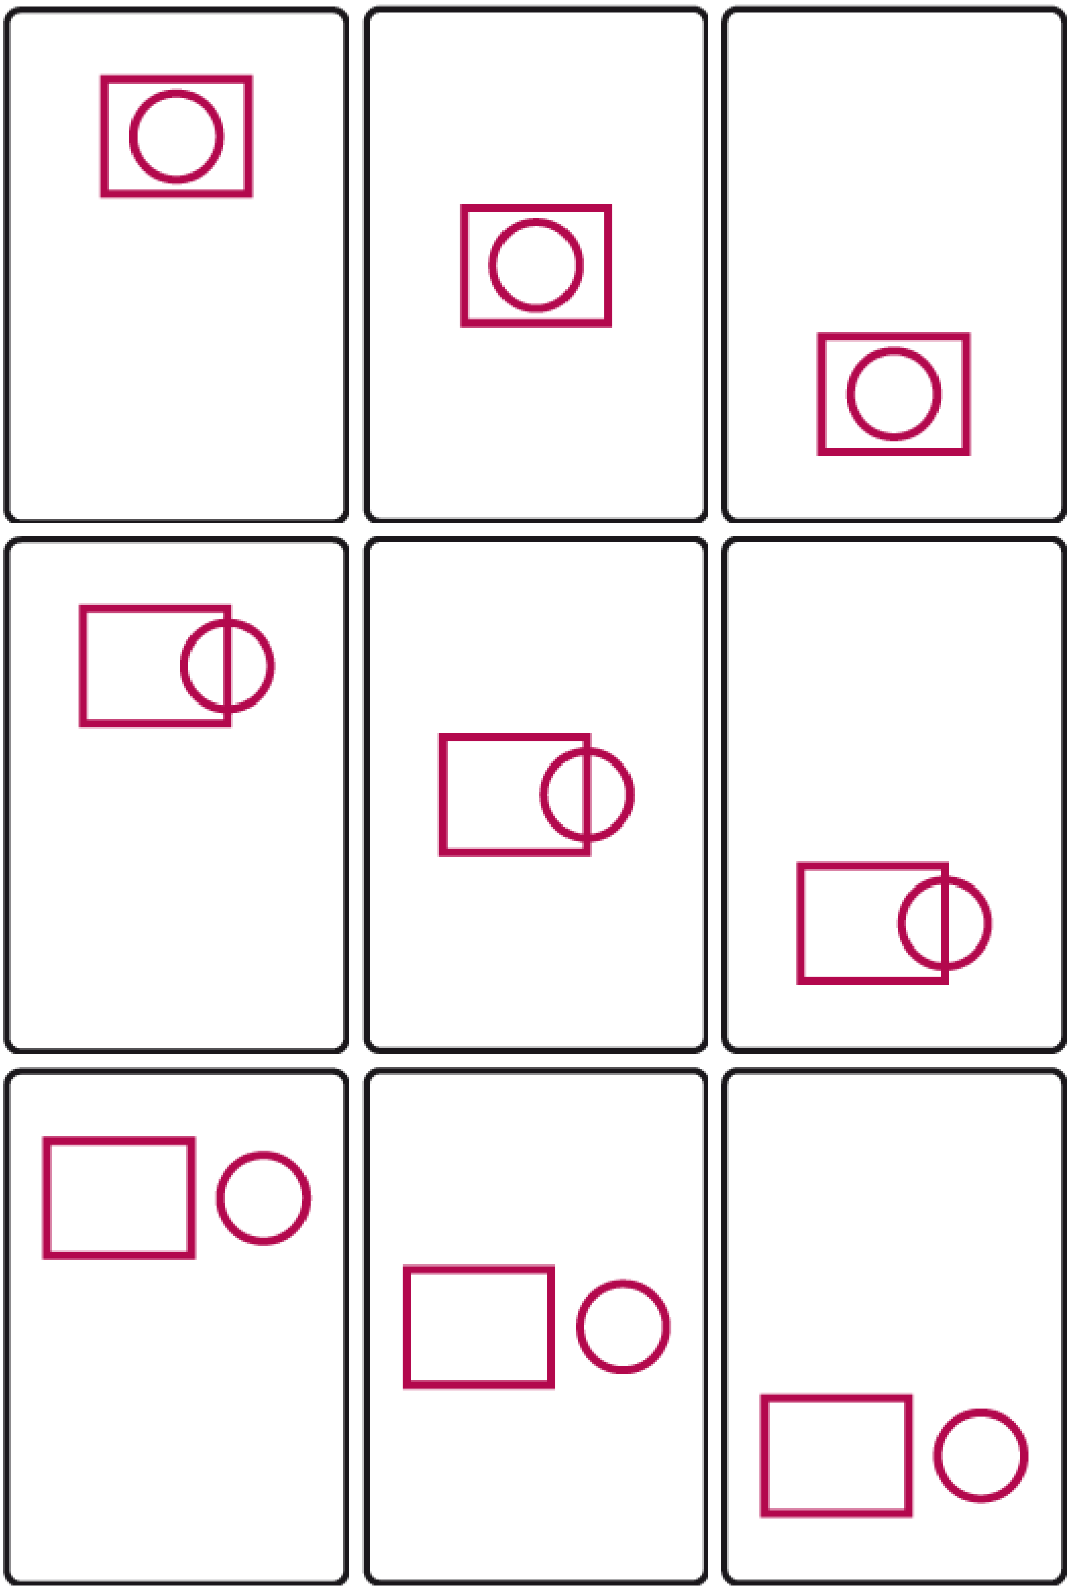

Appendix B. Card sorting experiment all stimuli

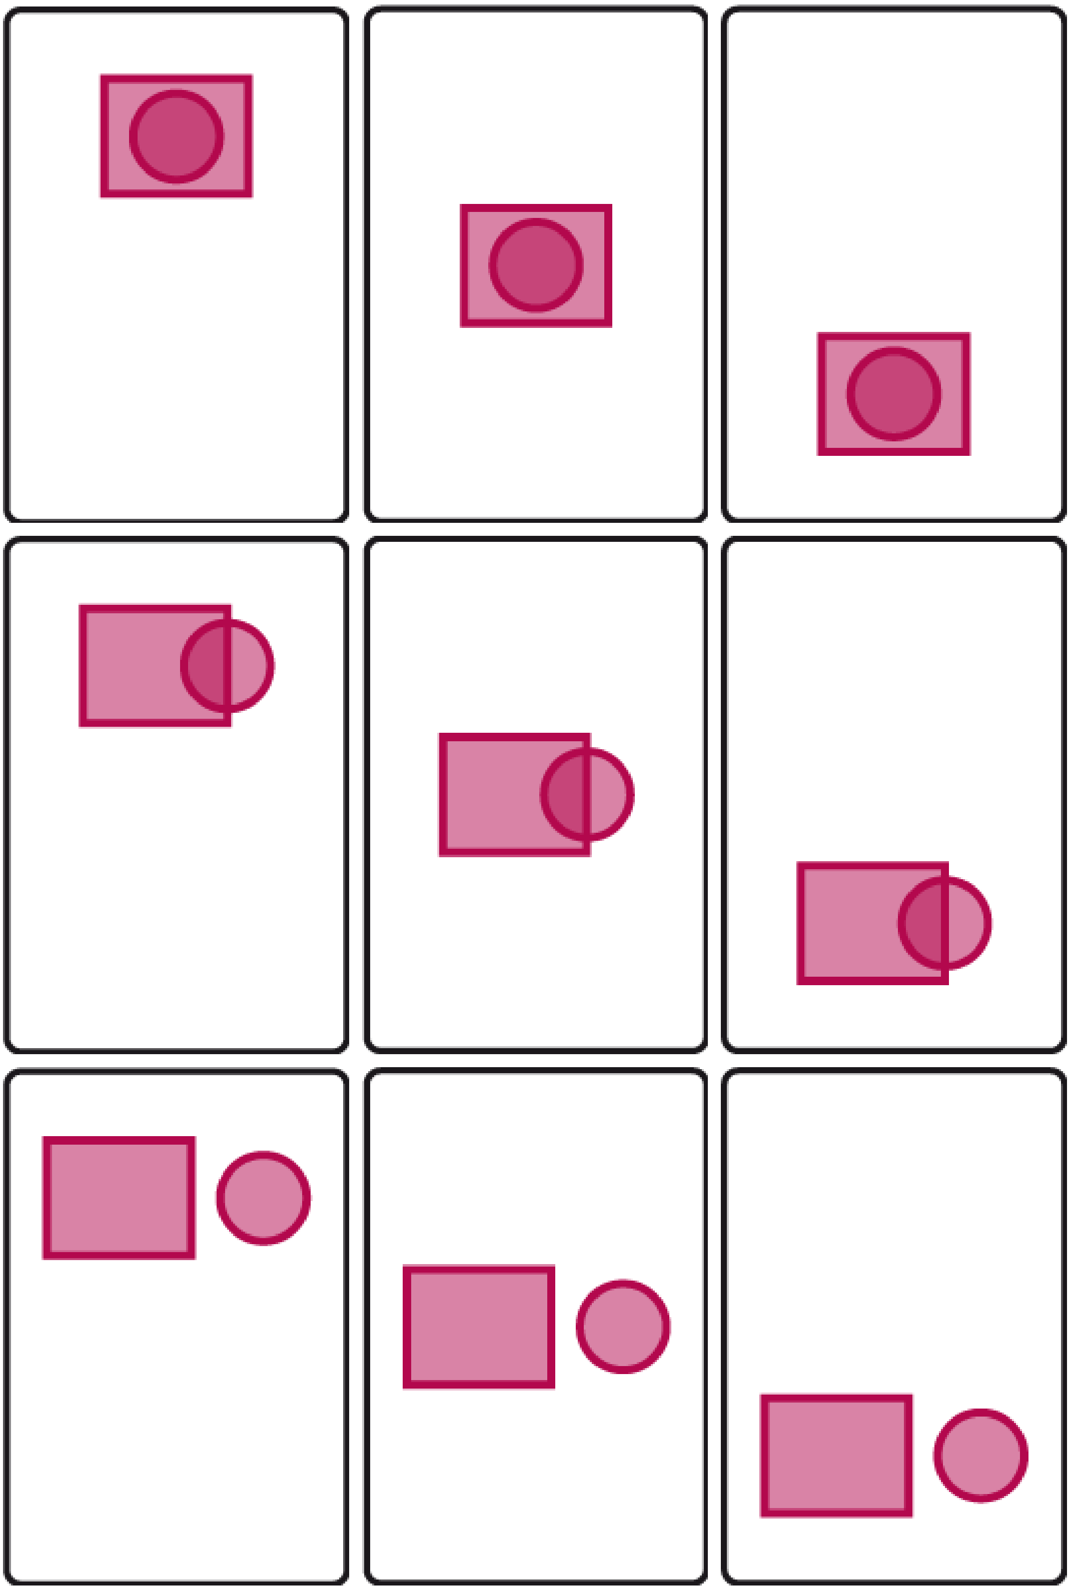

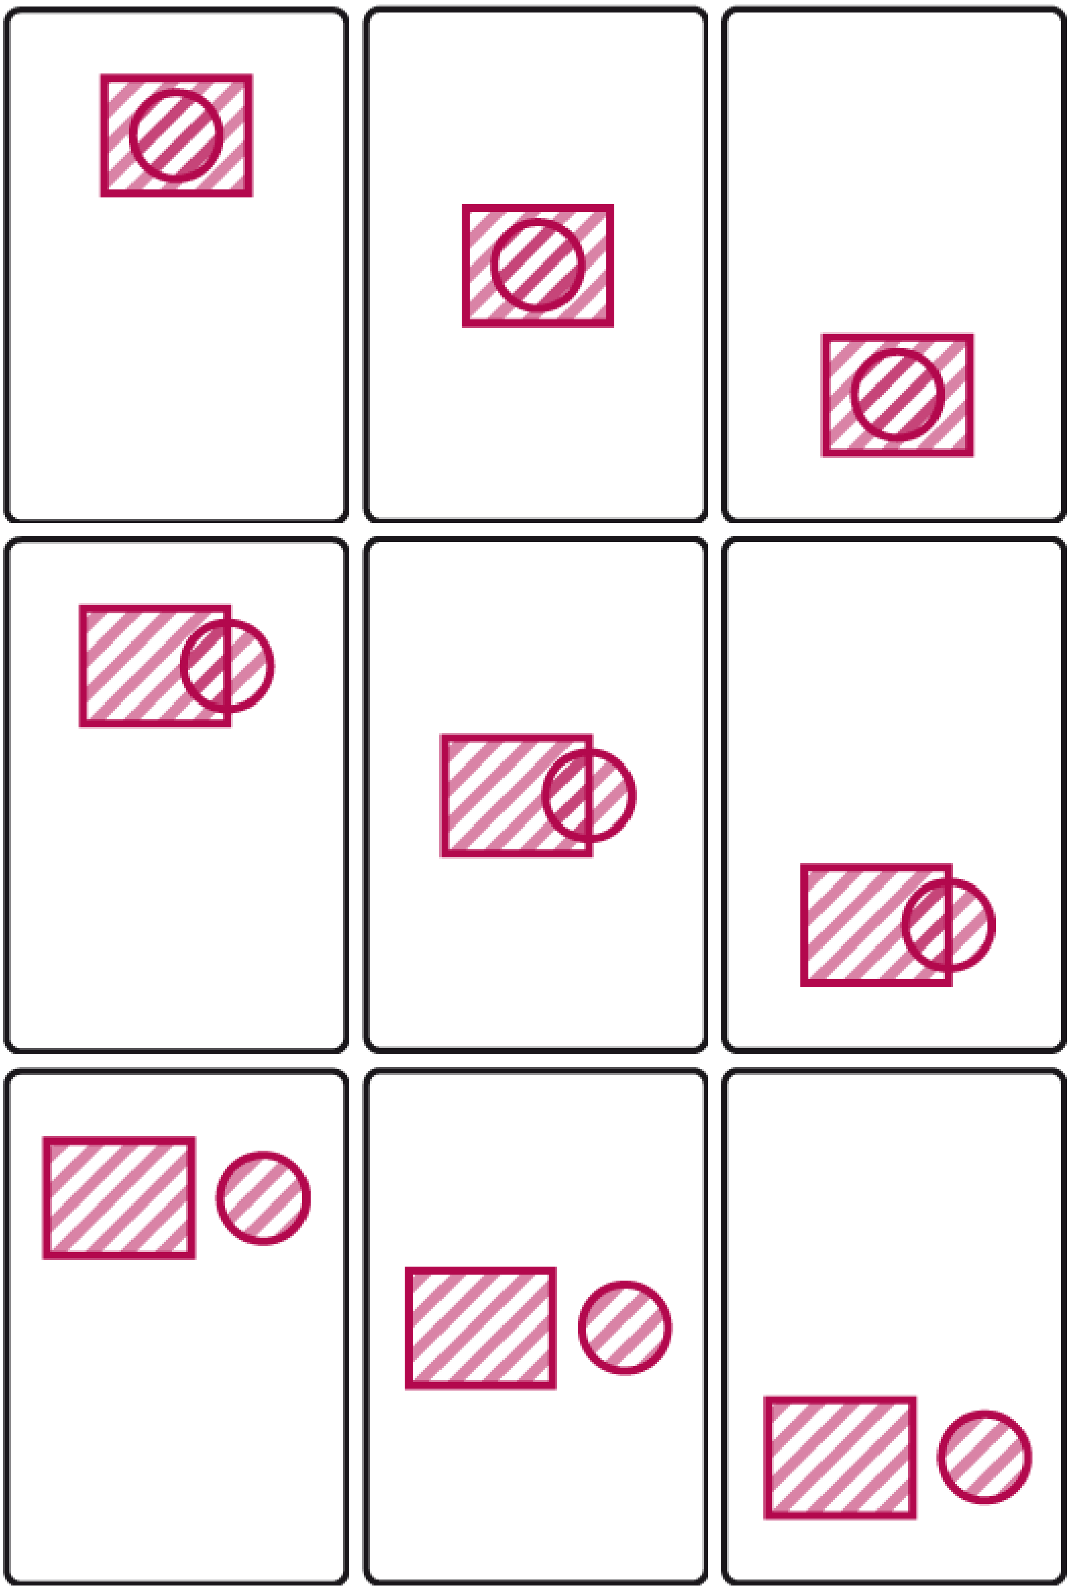

Appendix B. Card sorting experiment all stimuli

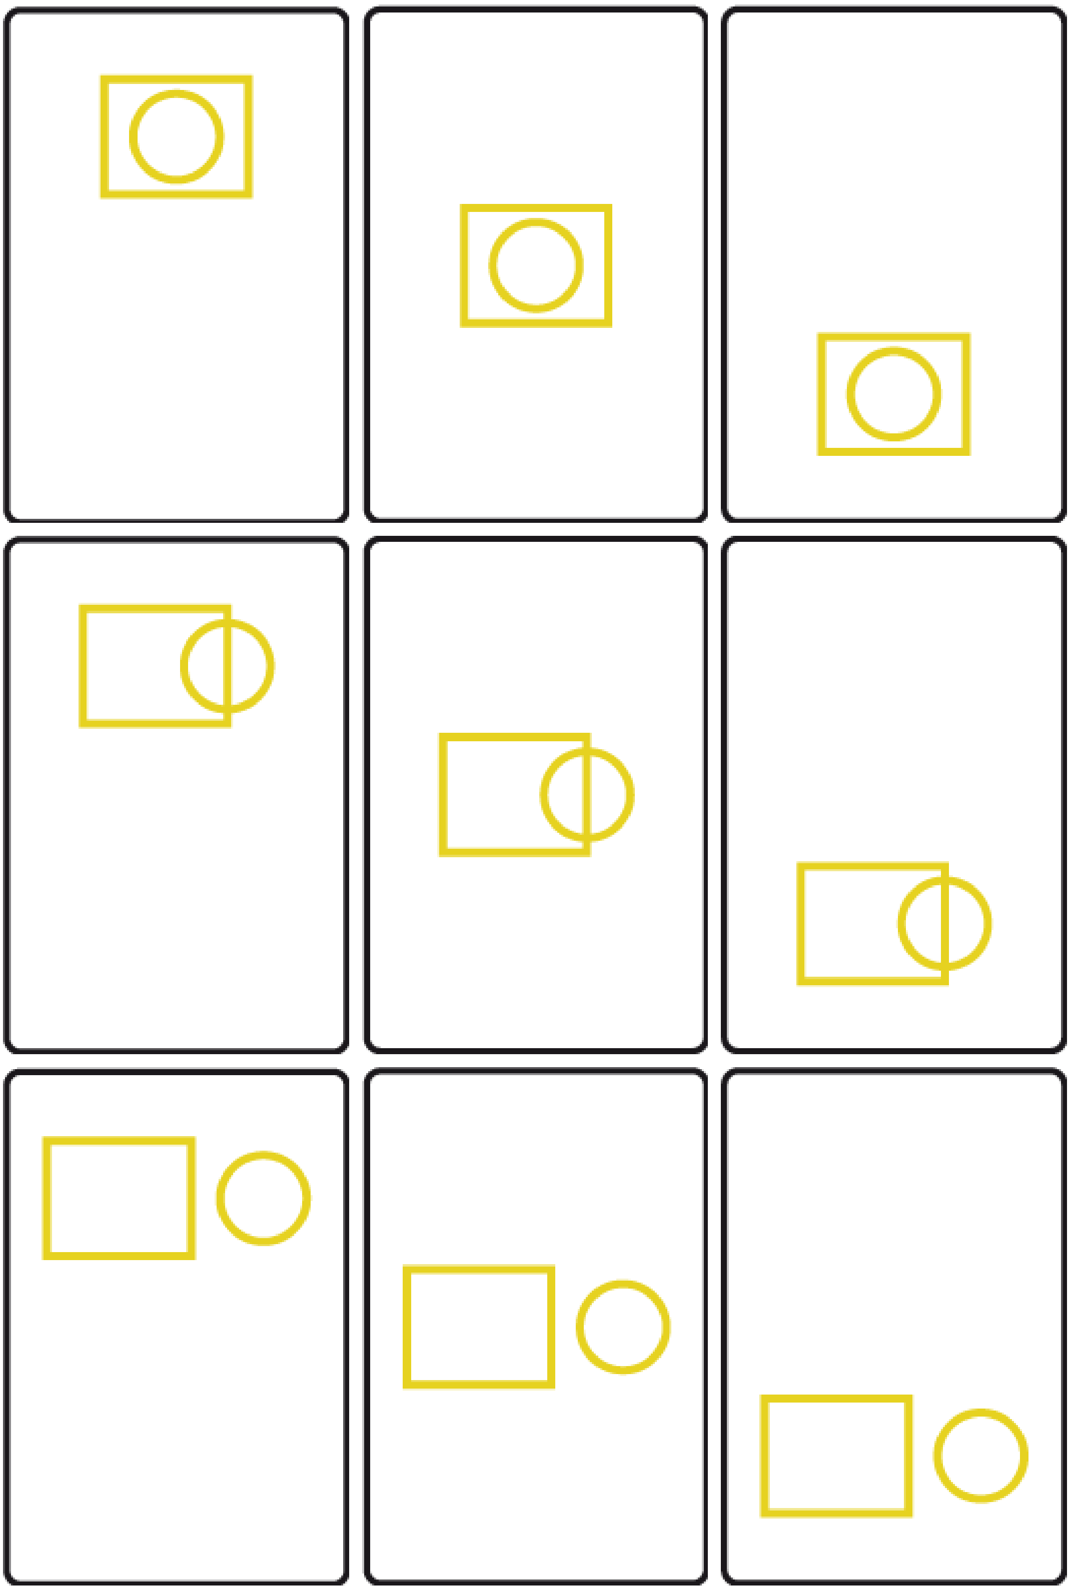

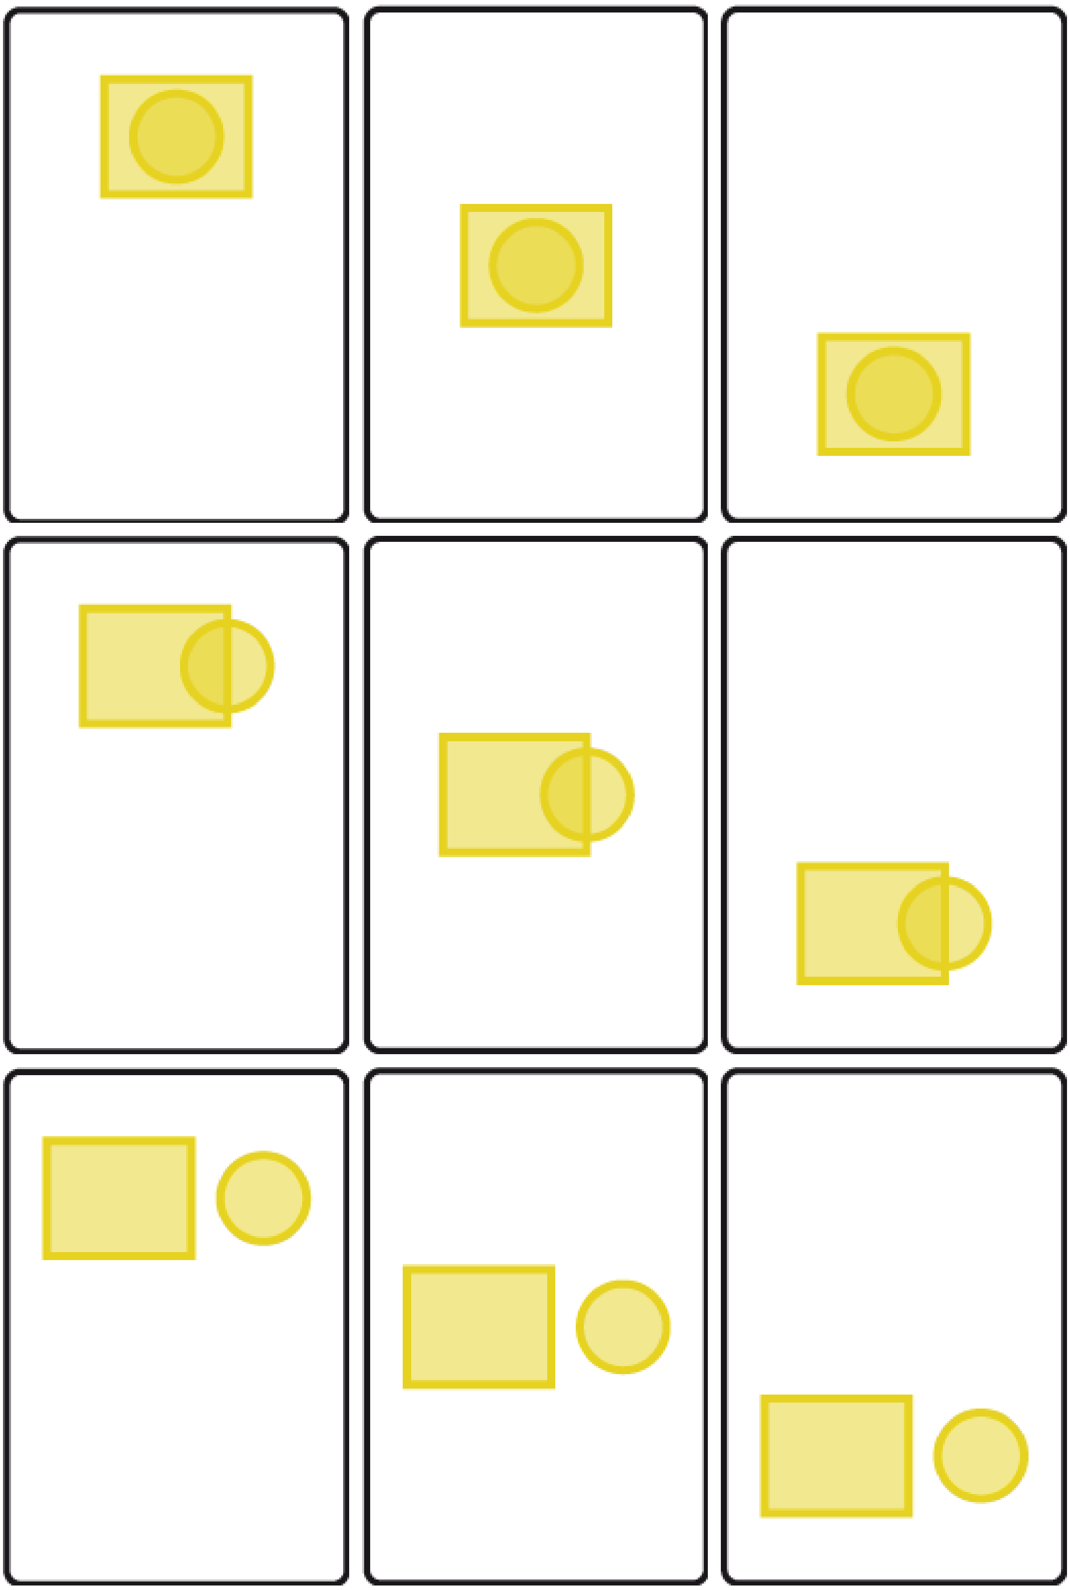

Appendix B. Card sorting experiment all stimuli