Strong Spatial Cognition

Christian Freksa

University of Bremen, Bremen, Germany
freksa@uni-bremen.de

Abstract. The ability to perform spatial tasks is crucial for everyday life and of great importance to cognitive agents such as humans, animals, and autonomous robots. Natural embodied and situated agents often solve spatial tasks without detailed knowledge about geometric, topological, or mechanical laws; they directly relate actions to effects enabled by spatio-temporal affordances in their bodies and their environments. Accordingly, we propose a cognitive processing paradigm that makes the spatio-temporal substrate an integral part of the problem-solving engine. We show how spatial and temporal structures in body and environment can support and replace reasoning effort in computational processes: physical manipulation and perception in spatial environments substitute formal computation, in this approach. The *strong spatial cognition* paradigm employs affordance-based object-level problem solving to complement knowledge-level computation. The paper presents proofs of concept by providing physical spatial solutions to familiar spatial problems for which no equivalent computational solutions are known.

Keywords. Cognitive architecture · spatial cognition · pervasive computing · ubiquitous computing · spatial affordance · knowledge representation

1   Introduction

The philosopher John Searle distinguishes "strong" AI from "weak" or "cautious" AI (Artificial Intelligence): "According to weak AI, the principal value of the computer in the study of the mind is that it gives us a very powerful tool. ... But according to strong AI, the computer is not merely a tool in the study of the mind; rather, the appropriately programmed computer really *is* a mind, in the sense that computers given the right programs can be literally said to *understand* and have other cognitive states. In strong AI, because the programmed computer has cognitive states, the programs are not merely tools that enable us to test psychological explanations; rather, the programs are themselves the explanations." [Searle 1980, p. 417]

Research in spatial cognition and the formalization of commonsense reasoning by means of logic has made substantial progress in reasoning about space and time in the past 25 years [Egenhofer and Franzosa 1991; Freksa 1991b; Cohn and Hazarika 2001; Renz and Nebel 2007; Ligozat 2011; Wolter and Wallgrün 2012; Dylla et al. 2013]. For example, approaches of qualitative spatial reasoning permit the computation of spatial relations that correspond to real or potential configurations in the physical

environment. These approaches employ AI tools (weak AI) to describe states of affairs in physical environments and to manipulate these descriptions in such a way that the resulting descriptions correspond to new states of affairs as obtained by physical operations in the spatial environment.

While weak AI models may achieve the same final configurations as the systems they model *(weak generative capacity, result equivalence, informational equivalence)*, only in certain cases they achieve these configurations in the same way *(strong generative capacity, strong equivalence* [Chomsky 1963], *computational equivalence* [Simon 1978]) as the corresponding operations in physical space. The reason is that only some spatial structures can be faithfully replicated in formal representations such as lists, trees, or arrays; other spatial structures and operations must be simulated in terms of the structures available in formal systems. As a consequence, computational operation sequences may be quite different from the sequences of mental and spatial operations modeled.

Whereas purely formal approaches employ sophisticated knowledge about a domain in order to infer new knowledge about spatial configurations, they may be only partially useful when it comes to modeling cognitive processes, their dynamics, their complexity, and their scalability [cf. Dreyfus 1979]. Here we require models that operate on the restricted domain knowledge of a cognitive agent and preserve the relevant structures of the system to be modeled; in spatial problem solving these are in particular spatial and temporal structures.

As spatial cognition does not exclusively take place in the mind – it also involves severe interactions between mind, body, and the spatial environment – a full model of spatial cognition must take the roles of the body and the spatial environment for spatial problem solving into account [see Chandrasekaran 2006]. In analogy to Searle's notion of *strong AI* we call embodied and situated models of spatial cognition that maintain the structural and functional properties of physical space *strong spatial cognition.* This paper expands on the work presented in [Freksa 2014; Freksa and Schultheis 2014; Freksa, to appear].

2   Real-World Spatial Problem Solving

We focus on a special class of real-world problems that are of particular significance for cognitive agents such as humans, animals, and autonomous robots: spatial and temporal problems in physical environments. These problems share basic structural properties that have been intensively studied in spatial cognition research [Freksa 1997] and are quite well understood today on the information processing level, i.e. the level that is directly accessible to computers and computer science. One outcome of this research is the insight that the best solutions to different types of spatio-temporal problems require a considerable variety of approaches and tools [Descartes 1637; Sloman 1985]. Some of the best approaches for human spatial problem solving make heavy use of the physical object level, i.e. they manipulate and perceptually inspect physical spatial configurations rather than solving problems entirely on the abstract information level.

In the following sections we will present a number of examples that illustrate different types of approaches to spatial problem solving.

Example 1: When a cognitive agent instantiates the route instruction *turn left at the next intersection*, he, she, or it does not require a detailed mental representation of the intersection or the environment; it also does not need to know and specify a precise turning angle before being able to follow the instruction.

The instruction provides a coarse guideline that permits the agent to move in an unfamiliar environment by perception-based route following and to select a *left* route from several alternatives that it may perceive in the vicinity of the intersection. The information about the turning angle is implicitly present in the spatial configuration that consists of the pose of the agent in relation to the route. The route instruction can be followed by means of a short perception-action loop. This does not require that the turning angle ever be made explicit in a cognitive representation.

In other situations cognitive agents may prefer to have detailed spatial knowledge before starting a spatial action as it may be easier to solve the problem by reasoning than by spatial interaction, as the following example illustrates:

Example 2: When I lost my keys that I last used during a trip some while ago, it may be worthwhile to reconstruct the preceding sequence of events on the trip mentally on the basis of remembered information about my preceding actions. Exploring the environment perceptually also might work to find my keys, but it could be more difficult or laborious, in the particular situation.

Embodied and situated cognitive agents are capable to operate on both the information processing and the physical object level. Perception and action operations serve as interfaces between the two levels; a memory serves to make information about the environment available to information processing in the absence of perceptual information and to store results of information processing for carrying out actions.

The classical cognitive science view treats cognition as a pure information processing activity [Simon 1978] that takes place entirely in the brain of a cognitive agent (respectively in a computer). But it was pointed out early on that the bodies of cognitive agents and the environments in which they perceive and act have significant effects on the types of solutions they pursue [Norman 1993; Clancey 1997; Wilson 2002; Wintermute and Laird 2008].

Furthermore, the information processing approach presupposes that a real-world problem has been comprehensively abstracted into an information processing task before cognitive processing can start. This assumption may be reasonable for routine tasks for which all necessary information is provided and which can be performed according to pre-existing standard patterns; however for novel problems, where a considerable part of the problem solving effort goes into identifying the information needed and finding appropriate representations, approaches, and tools, physical spatial configurations as in diagrams, maps, or other perceivable and / or manipulable

spatial configurations typically play an important role for solving problems [cf. Polya 1945]. In many cases, problem-specific approaches may be more appropriate and more efficient than general approaches. Specific approaches can take into account particular features of the problem domain to a larger extent than general approaches that abstract from specific characteristics.

In general, cognitive agents have a considerable variety of approaches and tools they can use to attempt solving a given spatial problem. In real-world cognitive problem solving we are confronted with problems where the identification of a suitable approach is more difficult than the computation of a solution on the basis of a given approach; once an appropriate approach has been selected, the problem solving procedure itself may be straightforward. This is illustrated in the following example:

Example 3: Suppose an agent's task is to determine visually (without a depth sensor) whether a tree is on this or on the other side of a fence (Fig. 1a).

A classical image analysis approach could use depth clues in the 2D projection of the 3D configuration to infer whether fence or tree is closer to the agent. Problem: the essential depth dimension is only poorly represented in the 2D projection. Spatial approach: Select a spatial reference frame that highlights the essential dimension; this can be achieved by relocating the agent such that the essential dimension is projected prominently onto the image of the configuration (Fig. 1b); now the task can be solved by considering only one image dimension, as the previous depth dimension has been mapped to the perceptually better accessible width dimension by spatial trans-formation in the problem domain.

| **Fig. 1a** Hard visuo-spatial decision problem | **Fig. 1b** Same problem in a more suitable spatial reference frame |

This specific example employs physical action only to modify perceptual acquisition of information from the environment without changing the 3-dimensional scene of interest. Note, however, that the 2-dimensional projection of the scene (that is typically the only information available for visual scene analysis) has considerably changed. Other kinds of manipulations would actually change the physical scene of interest in order to simplify or solve the spatial problem at hand.

The discussion shows that we need to address the question of *how to find* a suitable approach to solve a given spatial problem. Finding a suitable approach to solve a novel problem is one of the most interesting and challenging problems for cognitive agents. For the specific domain of spatio-temporal problems, we have good reasons to believe that the time is ripe to tackle this challenge; our belief is rooted in the fact that today we have a much better understanding of the properties of spatial and temporal relations and structures than twenty years ago.

There even is hope that once the challenges of identifying suitable spatio-temporal reasoning approaches to a given problem are tackled, we will be able to make use of the resulting approaches for addressing non-spatial problems, as well. This hope is rooted in the insight that human cognitive agents understand many problems through analogies [Gentner 1983] and metaphors [Lakoff and Johnson 1980]. Non-spatial problems may be solved by mapping them onto spatially constrained structures; here they may be solved more easily before mapping the solution back to the problem domain. This procedure would be in contrast to generalizing spatial approaches to unconstrained domains where we would employ highly general approaches.

The spatial (and to a lesser extent the temporal) domain is particularly accessible to autonomous mobile agents with visual, haptic, and auditory perception and memory, as well as with moving, turning, and grasping capabilities. These capabilities enable agents to flexibly interact with their environments. Specifically, agents can actively influence which parts and aspects of their environment they perceive and they can modify spatial configurations in the environment through their actions. So far, these capabilities have not been systematically investigated and exploited for cognitive systems architectures. We therefore propose to develop proof of concept implementations and demonstrations for solving spatio-temporal problems strategically by making use of spatio-temporal affordances.

A main motivation for studying physical operations and processes in spatial and temporal form in comparison to formal or computational structures is that spatial and temporal structures in the body and the environment can substantially support (and even replace) reasoning effort in computational processes [Dewdney 1988]. When we compare the use of different forms of representation [see Marr 1982], we observe that the processing structures of problem solving processes differ and facilitate different processing mechanisms [Sloman 1985]. Spatial structures that resemble the problem domain may result in a lower complexity than structurally deviating abstract representations, as they can make direct use of the inherent structural properties without a need for describing them [Nebel and Bürckert 1995].

A main objective of our work is to explore the scope of application of this principle. This will involve a representation-theoretic assessment of representational equivalence and similarity, on the level of both result equivalence and strong equivalence. We develop a framework to relate physical actions and perception activities [Bajcsy 1988; Lungarella et al. 2002] to information processing activities, in order to assess the trade-off between physical and mental operations. Such a framework has long been missing in the debate surrounding diagrammatic vs. analytic reasoning [Glasgow et al. 1995].

Our approach builds on well-established paradigms from cognitive science (e.g. *knowledge representation theory* [Palmer 1978], *affordances* [Gibson 1979], *knowledge in the world* [Norman 1980], *conceptual neighborhood* [Freksa 1991a]) and on research carried out in the collaborative research center SFB/TR 8 Spatial Cognition at the University of Bremen over the past twelve years.

3   Background and Motivation

AI research initially was concerned exclusively with mental aspects of cognitive systems, specifically with operations and processes that take place in the brain (respectively computer) [Feigenbaum and Feldman 1963]. Advances in robotics and knowledge representation have extended the scope of AI research to model perception and action processes, the (physical) bodies of agents, and the agents' spatial environments [e.g. Davis 1990]. The rather general structures of abstract formalisms used for knowledge representation in computers allow describing arbitrary aspects of bodies and environments in detail and to reason about them, including spatial and temporal aspects.

While abstract reasoning about the world can be considered the most advanced level of cognitive ability [see Freksa and Schultheis 2014], this ability requires a comprehensive understanding of mechanisms responsible for the behavior of bodies and environments. But many natural cognitive agents (including adults, children, and animals) lack a detailed understanding of their environments [Piaget 1929] and still are able to interact with them rather intelligently.

Example 4:   Children and dogs may be able to open and close doors in a goal-directed fashion without understanding the mechanisms of doors or locks on a functional level.

This suggests that knowledge-based reasoning is not the only way to implement problem solving in cognitive systems. Other systems of perceiving and moving goal-oriented autonomous agents have been proposed in biocybernetics and AI research to model aspects of cognitive agents [e.g. Braitenberg 1984; Brooks 1991; Pfeifer and Scheier 2001]. These models implement perceptual and cognitive mechanisms that follow physical laws rather than formal representations that follow the laws of logics. Such systems are capable of reacting to their environments intelligently without encoding knowledge about the mechanisms behind the actions and without the associated computational cost.

In our spatial cognition research we have investigated the potential of qualitative spatial relations, of structure-preserving schematic maps, and of the role of intrinsically spatial structures for spatial reasoning and spatial problem solving [Freksa 1991b, 2013; Schultheis et al. 2014]. A main result of this work is that structure-preserving representations can make direct use of spatial relations (e.g. spatial neighborhood, conceptual neighborhood, spatial order, and spatial orientation). Without structure-preservation, these relations would have to be derived through

knowledge-based processes in more abstract formal representations of space. Thus, spatial calculi that exploit structure-preserving representations can avoid the necessity of performing certain computational derivations, as they represent crucial relations intrinsically rather than extrinsically [Palmer 1978; Dirlich et al. 1983].

Spatial cognition research also has been concerned with issues of resolution and granularity, both on a physical and on a conceptual level [Hobbs 1985; Freksa and Barkowsky 1996; Mossakowski and Moratz 2012; Schultheis and Barkowsky 2011]. In knowledge representation, we must deal with the issue of level of detail on which we represent objects and configurations in order to solve certain problems. The finer the level of representation, the more problems we will be able to solve, in principle. But this comes at a cost: the more details we have to deal with, the more computation we have to invest. Cognitive processes frequently process information from coarse to fine rather than from fine to coarse. These processes are directly supported by physical and spatial properties of their environments. An illustration is given in the following example:

Example 5: In vision, *coarse* corresponds to *distant* and *fine* corresponds to *close*. The same sensor adapts its 'representation' of the world simply by physically moving towards an object or away from it.

The field of diagrammatic reasoning [Glasgow et al. 1995; Chandrasekaran 2006; Goel et al. 2010] is concerned with problem solving by means of diagrams, a special form of spatial representations. A key issue here is the comparison between formal and diagrammatic representations and reasoning processes. Of particular interest is the equivalence between the reasoning procedure operating on the corresponding formal structure and the problem solving procedure operating on the spatial structure. Strong equivalence has been mainly studied in comparing different formal systems. Comparing processes operating on physical spatial structures with processes operating on formal structures poses an interesting challenge, as we will require a reference framework that includes information processing and re-configuration of spatial configurations.

Our research builds on work in the areas of spatial and temporal reasoning and simulation, data structures, diagrammatic reasoning, mental representations, theories of knowledge representation and computation, and related areas. In the following I will sketch some of the issues that have been particularly relevant for our work.

In 1983, James Allen published a widely referenced paper on Maintaining knowledge about temporal intervals [Allen 1983]. In this paper the author developed a calculus for temporal relations based on the set of 13 jointly exhaustive and pairwise disjoint (JEPD) relations that can hold between two temporal intervals. Allen's approach became the role model for numerous calculi for qualitative spatial reasoning (QSR) [e.g. Guesgen 1989; Egenhofer and Franzosa 1991; Randell et al. 1992; Freksa 1992; Ligozat, 1993; Zimmermann, 1995; Moratz et al. 2000; Van de Weghe et al. 2005]. Whereas a single calculus is sufficient for reasoning about temporal relations, a multitude of calculi are required to cover all relevant aspects of spatial relations. In

particular, calculi for topological relations, for orientation relations, and for distance relations have been developed [Freksa and Röhrig, 1993; Cohn and Hazarika 2001].

Attempts to integrate the different aspects of spatial calculi in a single calculus failed. The reason for this is rooted in the fact that the different aspects of space are strongly intertwined as indicated by the following example:

Example 6: A topological relation between two objects constrains the distance between them; a set of distance relations between several objects constrains their relative orientations, etc.

An integrated calculus would have to compute and update all relations that are affected by a single change in order to maintain consistency among the relations. This would be computationally expensive and not useful if only a single aspect of the spatial configuration is of interest.

This computational dilemma points to a property of space that makes spatial structures particularly interesting: properties of spatial objects and configurations are intrinsically highly interdependent. If we modify one spatial aspect (e.g. distance, orientation, or topological relation) in a spatial structure, other spatial aspects will change 'automatically', as well. We call such a structure a *spatial substrate* [Freksa 2013]. If we move an object in space, the spatial locations of all its parts as well as their relations to other objects will change. If we change a single spatial aspect in a spatial substrate, all these changes take place ('for free'); no computing (or otherwise) effort is required.

In other words: The computational dilemma described above mutates into a special feature of spatial substrates: If suitable operations for spatial simulations in spatial substrates are available, we may be able to avoid an enormous amount of computation, whereby consistency is intrinsically guaranteed.

The use and exploitation of structural properties of representations makes up the core idea of data structures: a data structure supports a particular way of organizing information such that it can be used efficiently [Knuth 1997]. Depending on the problem structure and on the tasks to be performed on a representation, certain data structures may be better suited than others. Some data structures share structural properties with spatial substrates, in particular lists, trees, and arrays.

Example 7: Binary trees are particularly useful data structures for sorting and searching linearly ordered information, as the linear order can be mapped to the leaves of the tree while the branching structure of the tree corresponds to decisions to be taken in the sorting / searching process.

For other data, in which no order is implied in their appearance, an ordered data structure may be detrimental to the task as the structure may impose an unintended bias; therefore we employ structures that avoid such a bias, in these cases.

Example 8: Hash tables avoid a correspondence with spatial substrates and processes and thus permit data access independent of an ordering.

Unfortunately, we do not have suitable data and access structures for all operations that we would like to perform on computational representations of space, as the following example suggests:

Example 9: Zooming and perspective transformation operations involve computation on all elements of the domain and tend to be computationally expensive.

In spatial substrates, in contrast, we have more flexibility in accessing data. We can perform zooming and perspective transformation by manipulating the data access (perception) apparatus while leaving the data itself untouched. In the performance assessment of computer algorithms, data access operations usually are considered cheap in comparison to computational transformations.

Efficient use of information also depends on the operations we permit on data and knowledge structures. Unrestricted relations and operations may result in high complexity and in unfavorable scaling behavior while having no particular advantages, in many cases.

Example 10: In qualitative reasoning, we can relate each spatial (or temporal) relation to each other in the set of JEPD relations. When reasoning about a given domain, this leads to exponential complexity in the number of objects related to one another.

On the other hand, we know that in the spatial (respectively temporal) domain most transitions between relations cannot occur due to substrate-inherent reasons. This insight allows us to restrict a calculus to take into account only transitions between conceptually neighboring relations [Freksa 1991a]. This restriction on the representational level has no negative implications on the object level, as other transitions cannot occur in the represented domain. But it has great advantages: the calculus becomes tractable, as it results in only polynomial complexity [Nebel and Bürckert 1995].

Although spatial calculi have been applied with considerable success to a number of spatial problems [e.g. Wolter et al. 2008; Kreutzmann et al. 2013; Falomir et al. 2013; Dubba et al. 2015], there are at least two aspects of such calculi that seem in need of improvement. First, calculi mainly represent information about an abstract spatial problem, i.e., they largely fail to systematically exploit the constraints and affordances provided by the physical structure of the domain. As a result, spatial problems as represented by calculi may easily become computationally too complex to be solved efficiently, if spatial constraints (or subsets thereof) are not introduced on top of the spatial calculi – for example in the form of conceptual neighborhood (Freksa, 1991a; Nebel and Bürckert, 1995; Balbiani et al., 2000]. Second, there is currently only a poor understanding concerning which calculi are best suited for solving a given spatial problem: When faced with a specific problem it is not clear how to select among the many available formalisms for solving the problem.

In AI, early attempts to exploit spatial structures for reasoning purposes are found in the subfield of diagrammatic reasoning [Larkin and Simon 1987; Glasgow et al. 1995]. The idea was to make use of the spatial arrangement of objects on a (simulated) two-dimensional medium to optimize search and decision processes in reasoning. This idea was inspired by the way humans utilize diagrams employing their visual perceptual capabilities. Consequently, applications are found in geometric reasoning, in (physical) problem solving, and in the simulation of physical motion.

From a basic research perspective in cognitive science, spatial substrates play an important role in mental spatial reasoning using visual mental images [Kosslyn 1980, 1994] and spatial mental models [Johnson-Laird 1983, 1995]. In these mental representation formats, entities under consideration are dealt with in a similar way as they would be perceived and interacted with in the real world. The properties of these types of mental representations in mental spatial reasoning have been investigated and modeled from a computer science perspective by Schultheis and Barkowsky [2011] and Schultheis et al. [2014], among others.

In general, when a cognitive information processing system is analyzed from an informatics perspective, this can be addressed on three distinct levels [Marr 1982]: as a computational theory; from the perspective of knowledge representations and the processes operating on them; and with respect to a specific hardware implementation. It is an essential feature of our approach that we operate on all three levels: we are interested in spatial substrates that are physically realized and utilized by specific hardware; we operate on those substrates employing suitable representation structures and cognitive processes; and we aim at establishing a new paradigm of theoretically investigating cognitive systems from an information processing point of view.

In summary, we can identify numerous results which indicate that adaptations to specific task requirements may be cognitively and computationally more adequate [Sloman 1985] than the previously pursued goal of a 'General Problem Solver' [Newell et al. 1959]. In particular, exploiting the restrictions of spatial and temporal substrates may have considerable advantages over employing general abstract approaches. This insight also leads us to question the appropriateness of employing highly expressive representation languages for reasoning about the severely constrained spatial and temporal domains.

4   Objectives of this Work

Cognitive agents such as humans, animals, and autonomous robots comprise brains (respectively computers) connected by powerful interfaces to the environment: sensors and actuators. The sensors and actuators are arranged in their (species-specific) bodies to interact with their (species-typical) environments. All of these components need to be well tuned to one another to function in a fully effective manner. For this reason, we view the entire aggregate (cognitive agent including body and environment [Wilson 2002]) as a *full cognitive system* (see Fig. 2). This is in contrast to classical AI systems, which have focused on the structures and processes within the confinements of the computer.

Our research is concerned with the investigation of cognitive principles that govern the interaction between these high-level cognitive system components. Although the research is motivated by the capabilities we observe in natural cognitive systems, our goal is not to replicate or characterize a particular natural system. The project aims at investigating and analyzing the distribution, coordination, and execution of tasks among the components of embodied and situated spatial cognitive agents on the system level from a systems engineering / system analysis perspective [Pylyshyn 1988].
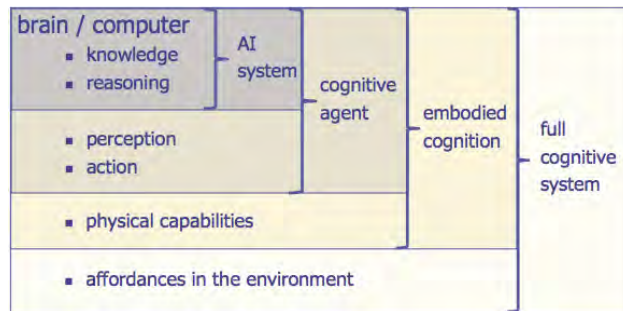
Fig. 2. Structure of a full cognitive system

In a classical information processing / artificial intelligence approach, we would describe the relevant components outside the brain or computer formally in some knowledge representation language or scheme in order to allow the computer to perform formal reasoning or other computational processing [Russell and Norvig 2010]. Physical, topological, and geometric relations would be transformed into abstract information about these relations and the tasks would be performed entirely on the information processing level, where physical, topological, and geometric relations are replaced by descriptions of their properties.

As we are dealing with spatial problems that originate from a spatial substrate and whose solutions eventually are relevant in a spatial substrate (e.g. we solve wayfinding tasks in order to apply the identified route in physical space), we will investigate under which conditions and to which extent we will be able to take advantage of the specific spatial properties and the structure of the problem domain. The goal is to use abstraction only in as far as it is useful for the given problem and to maintain the spatial structure when we can profit from its intrinsic spatial properties [Palmer 1978]. In this way we may be able to avoid effort, difficulties, and losses due to the problem abstraction process, the reasoning process (in a possibly not optimally adapted formalism), and the concretion process that maps the abstract problem solution back into the spatial problem domain.

When we talk of spatial structures and spatial reasoning in the context of spatial cognition, we implicitly include temporal structures and temporal reasoning, as we are concerned with cognitive processes and the dynamics of space which must take into account the structure and constraints of time.

Spatial and temporal structures are not as expressive as general abstract languages. Abstract languages can transcend the limitations of the physical realm. In comparison,

spatial and temporal structures have the advantage of representing precisely those configurations that we deal with in space and time. Limitations of spatial representation media have been explored in art (e.g. by Pablo Picasso and Maurits Cornelis Escher); their advantages come to bear when representing concrete spatial entities, events, or concepts that we imagine in terms of spatial relations or structures (e.g. abstract hierarchies). Here, the limitations of spatial structures are useful (e.g. when representing spatial configurations or abstract hierarchies on a spatial substrate like a piece of paper) as we do not have to make relations explicit that are implicitly provided by the spatial structure.

Thus, the question is not whether more general abstract representations or more specific concrete representations are better; rather, we need to consider in which form the problem is given and exactly which operations we want to perform. In the end, the question is how we can combine the advantages of a general abstract language with the advantages of a specific representation structure [Freksa and Barkowsky 1999].

As an example, consider geographic maps: they employ a spatial medium – paper or a 2D display of some sort. The spatial structure by itself is rather useless; a map requires symbols that add semantics and establish an abstract correspondence to entities in the real world. In principle, of course, everything that can be represented on a map could be described in terms of an expressive abstract language: all spatial relations that are implicitly expressed through the spatial medium could be made explicit in terms of linguistic or other descriptions.

For perceiving and acting cognitive agents who heavily rely on spatial interdependencies, such descriptions of maps would be difficult to use. Thus, we need to find the right balance between implicitly and explicitly available knowledge – or equivalently: between specific and general representation structures. In maps, this balance varies depending on the purpose of use: some features are expressed spatially (configurations, relative distances, and shapes) while others are expressed symbolically (type of road, type of land use) [Freksa 1999]. As both types of representation share the same spatial medium, they compete for kinds of interpretation, as the following example shows:

Example 11: In interpreting a road map, can I multiply the width of a road symbol on the map with the scaling factor of the map to determine the width of the corresponding road in the environment or is the width of the road symbol a constant related to that symbol?

The answer to this question depends on the design decision for the particular map type, namely which aspects of the environment are to be represented pictorially and which symbolically. In our research, we investigate distributions between implicit and explicit knowledge in intrinsic vs. extrinsic representations from a cognitive processing perspective: when we employ cognitive offloading of information processing from the mind (respectively computer) to the environment [Wilson 2002], we will have to add new information processing structures (on a more abstract level) to control the cognitive agent's use of the externalized knowledge. This creates

interesting trade-offs that we will investigate in the framework of the spatial substrate processing paradigm.

Just as the *logic programming* paradigm is designed to generate inferences about truth values by employing the laws of logics, we develop a spatial processing paradigm to generate inferences about spatial relations by employing the laws of space and time [Freksa and Schultheis 2014]. Logics is an excellent language on the meta-level, for describing states of affairs and for reasoning about them; our objective for developing the spatial processing paradigm is to produce an object level representation for processes of spatial cognition.

Against this background, our objective is to initiate a paradigm shift in representation and problem solving by placing emphasis on representational approaches and solutions that exploit knowledge in the world [Norman 1993] (and the affordances [Gibson 1979] and constraints [Freuder and Mackworth 1994] that come with it) as well as systematically investigating how to best distribute problem solving effort between abstract (knowledge *about* – meta-level) and concretely embodied (knowledge in the world – object level) modes of representation and processing. More specifically, our objectives comprise:

- Characterizing the division of labor between abstract knowledge representation and knowledge in the world.
- Devising representation and processing structures that facilitate the exploitation of knowledge in the world.
- Devising selection and control structures to identify a promising problem solving approach from a set of alternatives.
- Determining how (a) the effort required for building up an abstract representation; (b) the accessibility of knowledge in the world; (c) the frequency with which certain information is reused; and (d) the availability of computational and memory resources influence the division of labor.
- Developing measures that allow comparing effectiveness and efficiency of different combinations.
- Realizing the proposed approach in a simulation environment and on a robotics platform.

5   Approach

We focus on spatial and spatio-temporal tasks that are directly accessible by perception and allow for manipulation by physical action. This is the domain we understand best in terms of computational structures; we have well established and universally accepted reference systems to describe and compute spatial and temporal relations. The limitation to spatial tasks may turn out less severe than it may appear initially: numerous non-spatial problems can be transformed into equivalent spatial problems where the spatial structure helps to support the problem solving process. Human problem solvers make use of problem spatialization, for example, when visualizing a linguistically specified problem in the form of a diagram in order to better grasp the problem and/or to be better able to formalize it for formal problem

solving. Depending on the spatial representation chosen for the diagram, it may be easier or harder to grasp or formalize the problem.

The main hypothesis of the *strong spatial cognition* paradigm is that the 'intelligence' of cognitive systems is grounded not only in specific abstract problem solving approaches, but also – and perhaps more importantly – in the capability of recognizing characteristic problem structures and of selecting particularly promising problem solving approaches for given tasks. Formal representations generally do not facilitate the recognition of such structures due to a bias inherent in the abstraction. This is where *mild abstraction* can help as it abstracts only from few aspects while preserving important structural properties.

The insight that spatial relations and physical operations are strongly connected to cognitive processing will lead to a different division of labor between the perceptual, the representational, the computational, and the locomotive parts of cognitive interaction than the one we have been pursuing with AI systems: rather than putting all the 'intelligence' of the system into the computer, the proposed approach aims at putting more intelligence into the interactions between components and structures of a cognitive system as well as into the structure of the problem representation. More specifically, we aim at exploiting intrinsic structures of space and time to reduce the complexity of computation.

We argue that a flexible assignment of physical and computational resources for cognitive problem solving is closer to natural cognitive systems than the almost exclusively computational approach. For example, when we as cognitive agents search for certain objects in our environment, we have at least two different strategies at our disposal: we can represent the object in our mind and try to imagine and mentally reconstruct where it could or should be – this would correspond to the classical AI approach; or we can visually search for the object in our spatial environment. Which approach is better (or more promising) depends on a variety of factors including memory and physical effort required. Frequently a clever combination of both approaches will be best.

We develop and implement a proof of concept for the proposed approach to spatial problem solving through simulations of the perception and manipulation processes as well as through physical agent models, e.g. as generated by a 3D printer [Freksa 2013]. The research is primarily conceived as basic research in cognitive systems engineering: we identify and relate an inventory of cognitive principles and ways of combining them to obtain cognitive performance in spatio-temporal domains.

For this project, we can build on extensive research on spatial and temporal relations, their representation in memory, and with qualitative spatial reasoning in the framework of international interdisciplinary spatial cognition research. Naturally, the proposed approach will not be as broadly applicable as some of the approaches we have pursued in classical AI research, as it is intentionally restricted to spatio-temporal structures; but the approach promises to discover broadly applicable cognitive engineering principles for the design of tomorrow's intelligent agents.

Our philosophy is to understand and exploit pertinent features of space and time as modality-specific properties of cognitive systems that enable powerful specialized approaches in the specific domain of space and time. Since space and time are most

basic for perception and action and ubiquitous in cognitive processing, we believe that understanding and utilizing their specific structures will be particularly beneficial.

There are at least two general approaches towards studying cognitive systems: (1) by analysis or (2) by synthesis. A standard empirical approach would be to analyze an existing system in order to understand its functionality. When we are dealing with complex systems whose components interact in multiple ways, it becomes difficult to derive a single theory that explains the functionality of the underlying system.

In our research, we investigate cognitive systems by synthesizing components whose functions we understand. The objective is to provide a proof of concept in order to discuss and compare various system architectures. Our method follows the approach that Braitenberg called 'experiments in synthetic psychology'. As Braitenberg [1984] argues in his book *Vehicles*, this approach may lead in a straightforward way to a well-understood model that can be scrutinized by empirical methods.

Further, we can distinguish at least two types of models of systems that may or may not pursue different objectives: (1) models that aim at reconstructing the functional components, relationships, and performance of a system on a given level of abstraction and granularity [Zadeh 1979; Hobbs 1985] as closely as possible through replicating properties and functionality of their components (object-level models); and (2) models that make the scientists' knowledge about properties, relations, and functions of the modeled system explicit through descriptions or prescriptions (knowledge-level models). Both types of models are suited to enhance our understanding of cognitive systems and both have advantages and disadvantages.

Example 12: Well-known examples of the two types of models from a non-cognitive domain are wind channel models of automobiles or airplane wings (object level) and finite-element algorithms to power virtual wind tunnel simulations (knowledge level).

Before the theory of aerodynamics was mature enough to convey all relevant parameters that influence aerodynamic properties and computers powerful enough to simulate the aerodynamic interactions, engineers placed physical models (full sized or spatially scaled) that maintained crucial features such as shape and surface finish into the air flow generated by physical wind channels. In these wind channels, the engineers could measure physical forces to evaluate their designs under varying environmental conditions. As the theory of aerodynamics advanced (partly due to empirical testing of various aerodynamic shapes), the physical interactions between design and surrounding airflow were better understood; they could be characterized by finite element models, a numerical approximation approach to describing physical field behavior. Supercomputers were employed to calculate aerodynamic properties of cars and airplane wings, and other objects, as the computation of the interactions between all the parameters involves massive computation.

While the finite-element simulation has clear advantages over the wind tunnel simulation, it also has decisive disadvantages. Advantages are: the finite-element model reflects scientific understanding of aerodynamic interactions; furthermore wind

tunnel experiments required huge labs with special equipment that required personnel and consumed considerable amounts of energy. Disadvantages are: the mathematical simulation of the aerodynamic processes computes the effects of physical interaction by iterative numerical approximation processes; these processes have no direct correspondence to the physical interaction processes they represent and are not performed in real time. Considerable computation is required to integrate the results from a multitude of interactions.

For building cars and airplane wings the disadvantages of computational simulation are not significant; the simulations can run 'off-line' to compute the required characteristics of the design. In cognitive modeling, however, we are not only interested in the result of applying the model; we are interested in the dynamics of the cognitive processes themselves. Thus we may benefit from an object-level model that intrinsically guarantees certain domain properties – until we understand their significance sufficiently well to simulate them in a cognitive process model in real time.

In modeling computational problem solving in spatial substrates, we are confronted with the spatial substrate (object level) and at least two levels of abstraction: the knowledge level and a meta-knowledge level (Fig. 3).
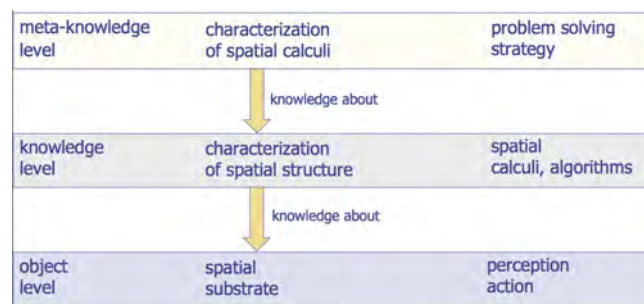

Fig. 3. Levels of cognition in spatial problem solving

The knowledge level makes aspects of the spatial structure on the object level explicit in the form of spatial relations and calculi. The meta-knowledge level makes knowledge about the relations and calculi on the knowledge level explicit in the form of knowledge about their use.

A fundamental contribution of computer science and artificial intelligence is to enable algorithmic interpretation of formally described knowledge such that formal representations of knowledge can be executed and thus perform as a model. In this way, the knowledge-level model can turn into a performance model whose behavior resembles that of the object-level system. To what extent it will be possible to reproduce the behavior of the object level system will depend on the structures of the representation and of the interpretation processes. It is not necessarily the most general or most powerful interpretation mechanism that will yield the best correspondence; rather, the best adapted knowledge structures and interpreters will win.

Object-level models do not make knowledge explicit; they maintain system properties and relations implicitly in their system structure. If the system structure and the process structure of the model closely match the corresponding structures of the cognitive system on the object level, its behavior can be expected to closely resemble that of the modeled system. Typically, the structures and processes of the modeled system are only partially known; the model designer fills in other parts on the basis of 'informed speculation'. Matching behavior provides no proof but may provide strongly suggestive evidence for the appropriateness of the model. Running and testing such object-level models and observing their global and detailed behavior can provide very useful information for further theoretical and / or empirical exploration of the cognitive system of interest.

Knowledge-level modeling can be considered the more sophisticated of the two methods, as it starts with knowledge that can be used for formal reasoning about scientific findings along established routes. A weakness of the knowledge-based approach, however, is that it forces the structure of the formalism and the reasoning procedures onto the system under investigation. This may not be a problem as long as we are only interested in the (static) initial and final states of problem solving processes of the cognitive system under investigation; however, if we are interested in modeling the dynamics of a cognitive system, a good match of module and process structures between model and target system become essential. This is where object-level models may score; they start with an engineering approach that initially focuses on system architecture and behavior. From here, valuable scientific knowledge on principles of cognitive processing can be derived. Braitenberg [1984] refers to this as the "law of uphill analysis and downhill synthesis".

In Section 2 I presented an example where spatial problem solving is supported by a suitable reference frame for a problem given in a spatial substrate (Example 3). In this example, the objects in the spatial substrate were not manipulated; only the reference frame was changed. In the following, I will present an example of how intrinsic properties of a spatial substrate can be exploited for spatial problem solving by manipulating spatial configurations. Manipulation for spatial problem solving constitutes a more severe interaction with the environment. In Example 13, the spatial problem configuration itself is modified in order to obtain a configuration that is equivalent with respect to the problem to be solved, but easier to analyze by spatial procedures.

Example 13: Suppose an agent's task is to identify the shortest route that connects a location A with a location B given several possible paths that can be chosen.

A classical knowledge-level approach would represent the lengths of the route sections, compute various alternatives of configuring these sections to connect A and B and determine the option with the smallest overall length. Observation: the lengths of the route sections need to be known and several alternatives have to be computed and compared before the one route of interest is identified.

Spatial approach [Dewdney 1988]: Here we use a mildly abstracted version of the street network: a paper map in which all paper regions which do not correspond to routes are missing. We obtain a deformable map consisting only of route representations which preserve the relative lengths of the route sections (Fig. 4a).

The map permits certain spatial reconfigurations of the network through deformation while preserving topology and important geometric constraints; in particular, an agent can (carefully) pull apart the positions A' and B' on the map (Fig. 4b) that correspond to locations A and B until a string of route sections forms a straight line between these positions (Fig. 4c); due to the geometric properties of the representation, the route sections corresponding to the sections on the straight line manifest the shortest route between A and B.

This approach avoids computation by reducing the problem to the relevant single dimension of length on which a basic geometric principle *straight line is shortest connection* can be directly applied. In both, Example 3 and Example 13, computational problem solving operations are augmented by spatial operations.
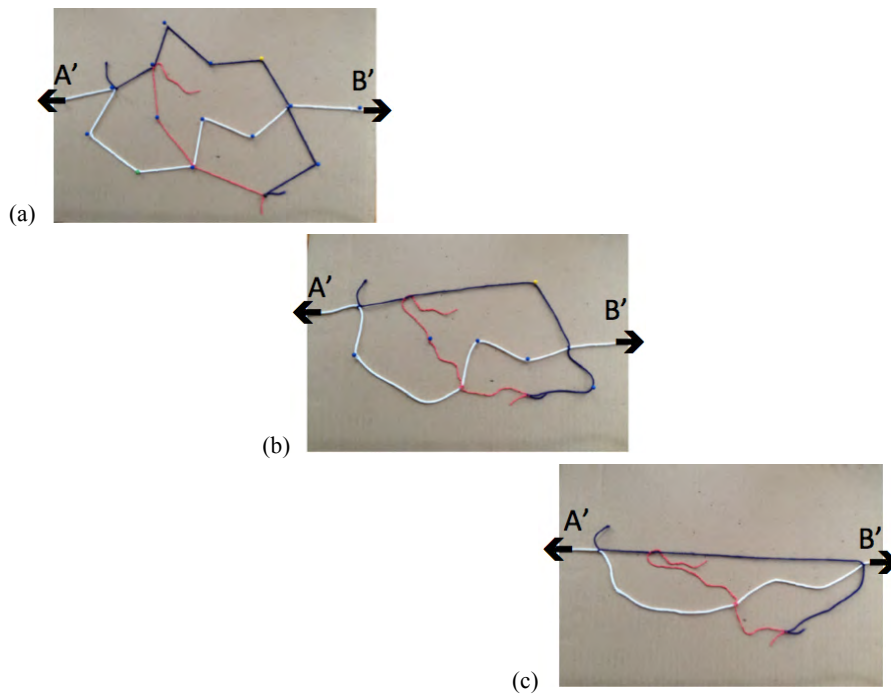

(a)

(b)

(c)

Fig. 4. Determining the shortest route from point A to point B by physical manipulation of a mildly abstracted representation of a route network. (a) The (non-elastic) strings corresponding to route segments preserve the relative distance relations of the original route segments; the distance relations are invariant wrt. physical manipulations (pulling apart strings at A' and B') which distort angles and shapes of the route network (b) and (c). The shortest route is identified as the route corresponding to the straight connection between A' and B' in (c).

6 Conclusions and Outlook

Our project sets out to investigate a new architecture of artificial cognitive systems that more closely resembles natural cognitive systems than purely knowledge-based AI approaches to cognitive processing. This is to be achieved by involving interaction with space through perception and action.

With today's availability of 3D printers, Example 13 can be implemented in the framework of a robotic system that interacts with and manipulates configurations in a spatial substrate: a route network can be extracted from an aerial photograph or from a map database and be printed (without the regions between the routes) on a 3D printer using a non-elastic and non-rigid printing substrate (comparable to paper). On the printer output, the robot's perception system identifies the start and end points on the route and grasps both points. The robot then cautiously pulls apart the two points until it can identify an almost straight connection between the start and end points of the route network; this connection will correspond to the shortest connecting route in the network. The example serves as a proof-of-concept for our project from which further explorations will follow.

The project brings together perspectives from a variety of disciplines: (1) the cognitive systems perspective, which addresses the cognitive architecture and trade-offs between properties of physical structures and properties of their descriptions; (2) the formal perspective, which characterizes and analyzes the resulting structures and operations; (3) the engineering perspective, which constructs and explores varieties of cognitive system configurations; and (4) the psychological-empirical perspective, which relates the effects of different system behaviors to those of natural agents. In the long term, we see potential technical applications of physically supported cognitive configurations, for example, in the development of future *intelligent materials* (e.g. 'smart skin' where spatially distributed computation is required that needs to be minimized with respect to computation cycles and energy consumption, and more robust and adaptable artificial agents, which can deal with unknown environments).

References

Allen, J.F.: Maintaining knowledge about temporal intervals, *Comm ACM* 26, 832-843 (1983)
Bajcsy, R.: Active perception, Proc. IEEE 76, 996-1005 (1988)

Balbiani, P., Condotta, J.F., Ligozat, G.: Reasoning about generalized intervals: Horn representability and tractability. TIME 2000: 23-29 (2000)

Braitenberg, V.: Vehicles: Experiments in synthetic psychology. MIT Press, Cambridge (1984)

Brooks, R.A.: Intelligence without representation, Artificial Intelligence 47, 139-159 (1991)

Chandrasekaran, B.: Multimodal cognitive architecture: Making perception more central to intelligent behavior, Proc AAAI, 1508-1512 (2006)

Chomsky, N.: Formal properties of grammar. In: Luce, R.D., Bush, R.R., Galanter, E. (eds.) Handbook of mathematical psychology II: 323-418. John Wiley and Sons, London (1963)

Clancey, W.: Situated cognition: on human knowledge and computer representations. Cambridge University Press, Cambridge, UK (1997)

Cohn, A.G., Hazarika, S.M.: Qualitative spatial representation and reasoning: An overview, Fundamenta informaticae 46 (1), 1-29 (2001)

Davis, E.: Representations of commonsense knowledge. Morgan Kaufmann, San Mateo (1990)

Descartes, R.: Discourse on the Method, Part VI. In: Newby, I., Newby, G. (prods.) The Project Gutenberg EBook 2008, #59 (1637)

Dewdney, A.K.: The armchair universe, W.H. Freeman & Company, San Francisco (1988)

Dirlich, G., Freksa, C., Furbach, U.: A central problem in representing human knowledge in artificial systems: The transformation of intrinsic into extrinsic representations, Proc. 5th Cognitive Science Conference, Rochester (1983)

Dreyfus, H.L.: What computers can't do. The limits of artificial intelligence. Revised edition. Harper & Row, New York (1979)

Dubba, K.S.R., Cohn, A.G., Hogg, D.C., Bhatt, M., Dylla, F.: Learning relational event models from video. J Artificial Intelligence Research 53, 41-90 (2015)

Dylla, F., Mossakowski, T., Schneider, T., Wolter, D.: Algebraic properties of qualitative spatio-temporal calculi. In: Tenbrink, T., Stell, J., Galton, A., Wood, Z.: Spatial Information Theory, COSIT 2013. LNCS 8116, 516-536. Springer, Heidelberg (2013)

Egenhofer, M., Franzosa, R.: Point-set topological spatial relations. Intern. J. Geographical Information Systems, 5, 2, 161-174 (1991)

Falomir, Z., Museros Cabedo, L., Castelló, V., González Abril, L.: Qualitative distances and qualitative image descriptions for representing indoor scenes in robotics. Pattern Recognition Letters, 34, 7, 731-743 (2013)

Feigenbaum, E., Feldman, J. Computers and thought. McGraw-Hill, New York (1963)

Freksa, C.: Conceptual neighborhood and its role in temporal and spatial reasoning. In: Singh, M., Travé-Massuyès, L. (eds.), Decision support systems and qualitative reasoning, 181-187. North-Holland, Amsterdam (1991a)

Freksa, C.: Qualitative spatial reasoning. In: Mark, D.M., Frank, A.U. (eds.) Cognitive and linguistic aspects of geographic space, 361-372. Kluwer, Dordrecht (1991b)

Freksa, C.: Using orientation information for qualitative spatial reasoning. In: Frank, A.U., Campari, I., Formentini, U. (eds) Theories and methods of spatio-temporal reasoning in geographic space. LNCS 639, 162-178. Springer, Berlin (1992)

Freksa, C.: Spatial and temporal structures in cognitive processes. In: Freksa, C., Jantzen, M., Valk, R. (eds.) Foundations of computer science. Potential – theory – cognition. LNCS 1337, 379-387. Springer, Berlin (1997)

Freksa, C.: Spatial aspects of task-specific wayfinding maps: A representation-theoretic perspective. In: Gero, J.S., Tversky, B. (eds.) Visual and Spatial Reasoning in Design, 15-32. Key Centre of Design Computing and Cognition, University of Sydney, Sydney (1999)

Freksa, C.: Spatial computing – How spatial structures replace computational effort. In: Raubal, M., Mark, D., Frank, A. (eds.) Cognitive and linguistic aspects of geographic space. Springer, Heidelberg (2013)

Freksa, C.: Strong spatial cognition (ext. abstract). In: Stewart, K., Pebesma, E., Navratil, G., Fogliaroni, P., Duckham, M. (eds.) Ext. Abstr. Proc. GIScience 2014. GEOinfo 40, 282-285, Vienna. Rev. version in Cognitive Processing 15 (Suppl 1): 103-105 (2014)

Freksa, C.: Computational problem solving in spatial substrates – A cognitive systems engineering approach. Int. J. Software and Informatics 9, 2, 2015 (to appear)

Freksa, C., Barkowsky, T.: On the relation between spatial concepts and geographic objects. In: Burrough, P., Frank, A. (eds.) Geographic objects with indeterminate boundaries, 109-121, Taylor and Francis, London (1996)

Freksa, C., Barkowsky, T.: On the duality and on the integration of propositional and spatial representations. In: Habel, C., Rickheit, G. (eds.) Mental models in discourse processing and reasoning, 195-212. Elsevier, Amsterdam (1999)

Freksa, C., Röhrig, R.: Dimensions of qualitative spatial reasoning. In: Piera Carreté, N., Singh, M.G. (eds.) Qualitative reasoning and decision technologies, 483-492. CIMNE, Barcelona (1993)

Freksa, C., Schultheis, H.: Three ways of using space. In: Montello, D.R., Grossner, K.E., Janelle, D.G. (eds.) Space in mind: Concepts for spatial education. MIT Press, Cambridge (2014)

Freuder, E., Mackworth, A. (eds.) Constraint-based reasoning. MIT Press (1994)

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

Gibson, J.J.: The Ecological Approach to Visual Perception. Lawrence Erlbaum Assoc., New Jersey (1979)

Glasgow, J., Narayanan, N.H., Chandrasekaran, B. (eds.) Diagrammatic reasoning: Cognitive and computational perspectives. AAAI Press, Menlo Park (1995)

Goel, A.K., Jamnik, M., Narayanan, N.H. (eds.) Diagrammatic representation and inference. Springer, Berlin (2010)

Guesgen, H.W.: Spatial reasoning based on Allen's temporal logic, TR-89-049, International Computer Science Institute, Berkeley (1989)

Hobbs, J.: Granularity. Intern. Joint Conf. Artificial Intelligence, 432-435 (1985)

Johnson-Laird, P.N.: Mental models. Harvard University Press, Cambridge (1983)

Johnson-Laird, P.N.: Mental models, deductive reasoning, and the brain. In: Gazzaniga, M.S. (ed.), The cognitive neurosciences, 65, 999-1008. MIT Press, Cambridge (1995)

Knuth, D.: The Art of Computer Programming, 1: Fundamental Algorithms (3rd ed.), Addison-Wesley, Reading (1997)

Kosslyn, S.M.: Image and mind. Harvard University Press, Cambridge (1980)

Kosslyn, S.M.: Image and brain - The resolution of the imagery debate. MIT Press, Cambridge (1994)

Kreutzmann, A., Wolter, D., Dylla, F., Lee, J.H.: Towards safe navigation by formalizing navigation rules. Intern. J. Marine Navigation and Safety of Sea Transportation, 7, 2, 161-168 (2013)

Lakoff, G., Johnson, M.: Metaphors we live by. University of Chicago Press, Chicago (1980)

Larkin, J.H., Simon, H.A.: Why a diagram is (sometimes) worth ten thousand words. Cognitive Science 11, 65-99 (1987)

Ligozat, G.: Qualitative triangulation for spatial reasoning. In: Campari, I., Frank, A.U. (eds.) Spatial Information Theory, COSIT 1993. LNCS 716, 54-68. Springer, Heidelberg (2013)

Ligozat, G.: Qualitative spatial and temporal reasoning. Wiley, London (2011)

Lungarella, M., Hafner, V., Pfeifer, R., Yokoi, H.: An artificial whisker sensor for robotics. IEEE Conf on Intelligent robots and systems (IROS), 2931-2936 (2002)

Marr, D.: Vision. MIT Press, Cambridge (1982)

Moratz, R., Renz, J., Wolter, D.: Qualitative spatial reasoning about line segments. In: Horn, W. (ed.) ECAI 2000, 234-238, IOS Press, Amsterdam (2000)

Mossakowski, T., Moratz, R.: Qualitative reasoning about relative direction of oriented points. *Artificial Intelligence* 180, 34-45 (2012)

Nebel, B., Bürckert, H.J.: Reasoning about temporal relations: A maximal tractable subclass of Allen's interval algebra. *JACM* 42 (1), 43–66 (1995)

Newell, A., Shaw, J.C., Simon, H.A.: Report on a general problem-solving program. Proc. Intern. Conf. Information Processing, 256-264. UNESCO, Paris (1959)

Norman, D.A.: The psychology of everyday things. Basic Books, Inc., New York (1980)

Norman, D.A.: Cognition in the head and in the world: An introduction to the special issue on situated action. Cognitive Science 17, 1-6 (1993)

Palmer, S.E.: Fundamental aspects of cognitive representation. In: Rosch, E., Lloyd, B.B. (eds.) Cognition and categorization, 259-303. Lawrence Erlbaum, Hillsdale (1978)

Pfeifer, R., Scheier, C.: Understanding intelligence, MIT Press, Cambridge (2001)

Piaget, J.: The child's conception of the world. Routlege & Kegan Paul Ltd., London (1929)

Polya, G.: How to solve it. Princeton University Press, Princeton (1945)

Pylyshyn, Z.: The role of architecture in theories of cognition. In: VanLehn, K. (ed.) Architectures for intelligence. Erlbaum, Hillsdale, N.J. (1988)

Randell, D.A., Cui, Z., Cohn, A.G.: A spatial logic based on regions and connection. KR 1992, 165-176 (1992)

Renz, J., Nebel, B.: Qualitative spatial reasoning using constraint calculi. In: Aiello, M., Pratt-Hartmann, I.E., van Benthem, J.F. (eds.) Handbook of Spatial Logics. Springer (2007)

Russell, S.J., Norvig, P.: Artificial intelligence: A modern approach (3rd ed.), Prentice Hall, Upper Saddle River, N.J. (2010)

Schultheis, H., Barkowsky, T.: Casimir: An architecture for mental spatial knowledge processing. Topics in Cognitive Science 3, 778-795 (2011)

Schultheis, H., Bertel, S., Barkowsky, T.: Modeling mental spatial reasoning about cardinal directions. Cognitive Science 38, 8, 1521-1561 (2014)

Searle, J.: Minds, Brains, and Programs. Behavioral and Brain Sciences 3, 417-457 (1980)

Simon, H.A.: On the forms of mental representation. In: Savage, W. (ed.) Perception and Cognition, 3-18. University of Minnesota Press, Minneapolis (1978)

Sloman, A.: Why we need many knowledge representation formalisms. In: Bramer, M. (ed.) Research and development in expert systems, 163-183. Cambridge University Press, New York (1985)

Van de Weghe, N., Kuijpers, P., Bogaert, B., Maeyer, P.: A qualitative trajectory calculus and the composition of its relations. First Intern Conf GeoSpatial Semantics (GeoS 2005), LNCS 3799, 181-211. Springer, Berlin (2005)

Wilson, M.: Six views of embedded cognition. Psychon. Bull. & Review 9 (4), 625-636 (2002)

Wintermute, S., Laird, J.E.: Bimodal spatial reasoning with continuous motion. Proc AAAI, 1331-1337 (2008)

Wolter, D., Dylla, F., Wölfl, S., Wallgrün, J.O., Frommberger, L., Nebel, B., Freksa, C.: SailAway: Spatial cognition in sea navigation. *Künstliche Intelligenz* 22, 1, 28-30 (2008)

Wolter, D., Wallgrün, J.O.: Qualitative spatial reasoning for applications: New challenges and the SparQ toolbox. In: Hazarika, S.M. (ed) Qualitative spatio-temporal representation and reasoning: Trends and future directions. IGI Global, 336-362 (2012)

Zadeh, L.A.: Fuzzy sets and information granularity. In: Gupta, M., Ragade, R., Yager, R. (eds.) Advances in Fuzzy Set Theory and Applications, 3-18 (1979)

Zimmermann, K.: Measuring without measures: The Δ-calculus. In: Frank, A., Kuhn, W. (eds.) Spatial Information Theory, COSIT 1995, LNCS 988, 59-67. Springer, Heidelberg (1995)