

Spatial Cognition

An AI Perspective

Christian Freksa*

Abstract. After a decade of temporal reasoning in Artificial Intelligence (AI) in the 1980s and 1990s, spatial reasoning and spatial cognition have moved into the focus of interest in concentrated research enterprises since the mid 1990s. This paper describes the interdisciplinary research area of spatial cognition from an artificial intelligence perspective and motivates the interest in the field and the challenges from a cognitive perspective. It argues that all themes of cognitive science surface in spatial cognition and that spatial cognition is particularly suitable to investigate these themes. The particular significance of spatial structures for knowledge acquisition and knowledge processing by cognitive agents is described; it is shown why spatial structures are instrumental in making sense of physical environments and abstract worlds. Basic approaches to computationally process spatial knowledge are sketched out; the role of qualitative reasoning in AI is compared to the role of qualitative approaches in other disciplines. The relative merits of intrinsically spatial and of more abstract, non-spatial ways of dealing with spatial knowledge are discussed. The role of schematic representation of spatial knowledge is addressed.

1 WHAT IS SPATIAL COGNITION?

Spatial Cognition is concerned with the acquisition, organization, utilization, and revision of knowledge about spatial environments. It allows cognitive agents to act and interact in space intelligently and to communicate about spatial environments in meaningful ways. Spatial and temporal cognitive capabilities enable humans to manage basic and high-level cognitive tasks in everyday life.

1.1 Natural and artificial, real and abstract systems

The cognitive agents that are investigated in spatial cognition are natural agents (humans and other animals); artificial agents (robots); they may exist in the real physical world; or they may be conceived as abstract notions or described by abstract formalisms (simulation programs).

The fact that we study such different systems in the common framework of spatial cognition underscores that we are looking for the underlying information processing principles shared by these different systems. In fact, each of the domains cannot be successfully investigated on its own, as we lack methods in each domain to discover all relevant features. For example, in natural systems, we can study behavior, anatomy, physiology, and we can perform brain imaging studies. However, sometimes we do not know local functions of the agent's neural system and we do not know

whether the agent of our investigation performs a spatial task (and if so, which one) in a given situation; we cannot validate certain experimental results, as the agents may behave differently due to learning or other internal or external factors beyond our control. Furthermore, different cognitive agents may apply different strategies to solve a given task; in such cases, we cannot use other cognitive agents to confirm our findings. Moreover, for functional and / or ethical reasons we are restricted in our investigation methods.

The situation is quite different in artificial systems: while we have precise knowledge of the local structures and functionalities, it is frequently difficult to design systems to produce a certain behavior; we may not know which representational and computational paradigm will be most appropriate and we may not even know the computational task to be performed to achieve a desired result. Similarly, investigations of real and abstract systems have different strengths and weaknesses which complement each other rather well. It is always helpful and usually most efficient if we can answer certain questions abstractly; on the other hand, investigations on real systems are required to ask the right questions.

While natural systems provide us with inspirations about purpose and function of cognitive systems, artificial structures allow us to investigate processes we cannot or do not want to study in living organisms and other integrated structures. Thus, we study natural and artificial systems both as real physical systems and as abstract entities to unveil the secrets of cognitive principles.

1.2 An interdisciplinary enterprise

Spatial cognition is a subfield of cognitive science. Almost all disciplines involved in the cognitive science enterprise contribute to spatial cognition: artificial intelligence, cognitive psychology, linguistics, philosophy of mind, psychophysics and cognitive neuroscience, cognitive anthropology, mathematics, cognitive geography and cartography, architecture and design. Each of these fields employs specific approaches to uncover certain aspects of cognitive structures and processes. For example, cognitive psychology studies natural cognitive systems and their behavior by conducting empirical investigations and analyzing their results; AI studies cognitive functions by constructing formal and synthetic structures and algorithms and empirically testing their behavior; and linguistics combines the analysis of natural language with the construction of formal structures to support or refute theories of spatial language. Geographical information science makes use of results from these disciplines [17].

The different approaches employed in cognitive science can be combined rather effectively to gain insight into complete cognitive mechanisms. Philosophers of mind and AI researchers investigate

* Cognitive Systems Research Group and SFB/TR 8 Spatial Cognition, Universität Bremen, Germany. Email: freksa@sfbr8.uni-bremen.de

conceptualizations of abstract and real spatial entities (e.g. borders between countries and borders between valleys and mountains) and study their formal properties. This provides a starting point for the construction of abstract formal systems according to criteria provided by empirical studies on real systems like autonomous robots. The implementation of abstract structures in real physical systems provides important insights into the flaws of the theory; jointly with empirical scientists AI researchers compare the performance of these artificial systems with the corresponding functions of the natural systems.

1.3 Uphill analysis and downhill invention

Cognitive systems tend to be complex. As a consequence, only partial models can be constructed from empirical data that reliably reflect the structures and functions of the natural role models. On the other hand, running computational models must be complete on a given level of description to be executable. Thus, functional AI models must contain structures that are not based on empirical evidence but on the constructor's intuitions. A great advantage of synthetic systems is that all constructive elements and structures are known, at least in principle.

If we view natural and artificial agents as different implementations of a given cognitive functionality, we can actively explore differences in performance and adapt our implementations according to new insights. In this way, cognitive psychologists may extend their range of empirical studies to artificial agents and AI programs can help bridge the knowledge gap between structure and function of cognitive systems.

The neuroanatomist and cybernetician Braitenberg characterized the use of synthetic constructs for the exploration of natural systems by noting that a given performance always can be achieved by many different mechanisms [5]. He formulated the *law of uphill analysis and downhill invention* noting that it is easy to create little machines that produce surprising behavior by simple means and much more difficult to derive from the outside the internal structure from the observation of behavior. As a psychological consequence of this he noted that we tend to overestimate the complexity of a mechanism when we analyze it. His experiments suggest that by building and exploring synthetic structures on the basis of biological principles we may make the best progress towards understanding natural structures.

2 WHAT IS SPECIAL ABOUT SPACE?

The discussion of the interdisciplinary methodology in the previous chapter applies to cognitive science in general. In the present chapter I will discuss the distinctiveness of *spatial* cognition.

2.1 Subjective perception and objective measurement

Our knowledge about physical space differs from other knowledge in a very significant way: we can perceive objects in space directly through various senses and modalities. Unlike for other perceivable domains, spatial knowledge obtained through one sensory channel can be confirmed or refuted by perception through other channels. Thus, even though perception is subjective, it may be reliable to a greater or a lesser degree.

The reliability of perceptual knowledge depends on a second factor: relative vs. absolute judgment. Relative judgments are much more reliable than absolute judgments. This fact is well-known in psychophysics and it is particularly true for spatial judgments. It is much easier to reliably compare two sizes, orientations, positions, etc. than to estimate sizes, orientations, or positions in absolute terms. The reason is that perception essentially works by comparison: if comparative information is not available in the perceptual domain, comparison must be made with internal (memorized) values; as this requires chains of comparisons or judgments, it is much more error-prone.

Reliability of knowledge also may be affected by directness of access. For example, knowledge about time cannot be obtained directly through perception; it requires memory and is derived on the basis of knowledge about space; knowledge not at all related to perceptual information may be even harder to acquire.

A great advantage of the spatial domain over other domains of cognition is that spatial distances, spatial orientations, and temporal durations in a physical environment can be determined objectively by means of technical measuring devices. Furthermore, we can derive spatial and temporal relations by means of a solid theory of geometry and by mathematical calculi. As we have an objectively accessible reference world as object of cognition, we can apply the methodology of the natural sciences and derive 'hard' results about cognitive processes similarly as physics derives laws of mechanics through measuring distances and durations in spatial environments.

2.2 Bridge between abstract notions and physical world

From a cognitive point of view, the spatial domain is special for an additional – though related – reason: through the fact that spatial knowledge is available to humans (and other cognitive agents) through various senses and modalities (e.g. visually: sizes, angles, textures, hue, saturation; acoustically: loudness, orientation, sound spectrum; haptically: size, shape; proprioceptively: orientation, distance) we are disproportionately confident about our knowledge about spatial relationships as compared to other features in our environment: we take them *for real*.

In addition to perceiving space multimodally, we can modify spatial configurations by moving or otherwise manipulating physical objects. Thus, we can expose the environment to the perceivable influence of physical forces and we can modify our spatial perception by moving ourselves. In this way, we experience the structures and laws of space as we can experience no other dimension. The spatial domain therefore serves as a particularly strong interface between the abstract world of cognition and the concrete physical environment.

From a developmental point of view, it appears plausible that spatial concepts are among the first to be formed in natural cognitive agents; basic spatial or spatio-temporal relationships can be established in a rather straightforward senso-motoric feedback loop and do not require language or an understanding of the world, as is impressively demonstrated in research on 'embodied cognition' [19, 2]. If spatial concepts and spatio-temporal relationships are among the first cognitive concepts formed, they become prime candidates for forming the basis for new concepts. In this way, more abstract concepts can be linked to more concrete concepts that are rooted in the real world through the perceptual interface.

3 SPACE-BASED CONCEPT REPRESENTATION

Cognitive representations can make use of perceptual structures that have been developed during the process of adaptation to the environment. In the space-based approach that will be outlined in this chapter, knowledge is represented in terms of environment-related concepts like *in front of* or *behind* that are anchored in spatial relations between physical objects in the real world, rather than in terms of general abstract notions like *zero* or *one*, *on* or *off*, *true* or *false*. It has been shown in the cognitive science literature that explicit semantic connections to the problem domain can be processed more appropriately than their logic abstractions [20].

3.1 Domain-specific vs. general representation format

What difference does the relation to the spatial environment make? First of all, we recognize the disadvantage of representations that are connected to specific entities: these representations are not as general as abstract representations that could be instantiated by arbitrary entities. This is a reason why we introduce the advantages of algebra over calculus in mathematics and the exploitation of tautologies over truth table manipulation logics.

As long as we deal with complete and certain knowledge, abstract representations in fact are superior as certain reasoning steps may be taken without reference to the object domain; if we are dealing with incomplete and uncertain knowledge, the situation may be different. In the following, I will consider the use of space-based representations to reason about spatial relations.

3.2 Exploiting spatial structures for reasoning

The power of abstract reasoning does not come for free. When cognitive agents must reason on the basis of incomplete or uncertain knowledge – the standard situation for cognitive agents – they may only be able to draw logic inferences if they can substitute for the missing knowledge by good guesses. Good guesses may come from a good theory or from well-informed knowledge sources. A good theory may provide default values that can support reasoning processes in standard situations. A well-informed knowledge source may complete missing knowledge just as needed. The former approach employs understanding of the general situation while the latter approach employs informants that need not understand what is going on.

3.2.1 World knowledge and domain knowledge

The interesting case is the case of the ‘good theory’. ‘General knowledge’ comes in two varieties: (1) relating to the whole, or every member of a class or category (the mathematical sense of ‘general’); and (2) being usually the case; true or applicable in most instances, but not all (the colloquial sense of ‘general’). Here again, the former is the more interesting: knowledge that always holds will be even more valuable for logic inferences than knowledge that usually holds. Can we separate the two types of knowledge?

Knowledge that always holds is world knowledge; it expresses principles or laws of the world. Knowledge that usually or frequently holds is domain knowledge. Both types of knowledge are useful, but if we can separate the two, we may be able to distinguish necessary and plausible consequences.

Spatial structures can provide us with knowledge about necessary implications of spatial relations and actions. They can be used to replicate spatial relations in the real world and thus implicitly fill in missing relations that we don’t explicitly know. Spatial and spatio-temporal structures have the advantage of being general enough to support a large scope of cognitively relevant features and phenomena and restrictive enough to provide useful constraints for reasoning. Our hypothesis is that spatio-temporal phenomena form the basis for our thoughts and reasoning and can be best supported by a spatio-temporal substrate.

3.2.2 Sources of uncertainty

When we deal with uncertain knowledge, it is crucial that we understand the nature of the uncertainty. Cognitive agents must live with only partial knowledge about their environment and still take the right actions. Lack of knowledge about precise locations, durations, or actions affects local uncertainty much more than global uncertainty. It is known that children learn laws of persistence and dynamics at an early age. Knowledge about spatio-temporal coherence allows cognitive agents to make reliable predictions despite missing knowledge [7]. In our everyday actions in traffic etc. we rely on spatio-temporal coherence to control our actions.

This insight will enable us to resolve the uncertainty appropriately. Too abstract representations of the environment may only tell us about alternative options in terms of logical disjunctions; this may be not enough to resolve uncertainty intelligently. In spatio-temporal worlds, uncertainty almost always is a function of spatio-temporal vicinity due to the relevant laws of physics; thus, when we employ spatio-temporal representations that maintain spatio-temporal coherence, we get information about vicinity for free. This information can be used to resolve perceptual and inferential uncertainty.

3.3 Reasoning about non-spatial relations

In the previous chapter I have given arguments, why space-based representations may be useful for reasoning about spatial situations. In the present chapter I will argue that space-based representations may support cognition beyond the spatial domain.

3.3.1 Concept acquisition through perceptual experience

Let us go back to the concept acquisition process I discussed in section 2.2. Suppose we have a cognitive agent equipped with spatial concepts but no other concepts about its environment. How can children (or adults or artificial cognitive agents) acquire and communicate concepts that can not be directly related to the physical environment through perception? Cognitive agents will be able to form new concepts related to the environment if (1) a perceptual, (2) a conceptual, or (3) a combination of a perceptual and a conceptual pattern gives rise to a new concept.

In other words, new concepts can be formed if they have something in common with existing concepts or perceptions. If they have everything in common with existing concepts, they are no new concepts. If they have nothing in common with existing concepts, they cannot be related to experiences and therefore they cannot be generated. Mathematically speaking, new concepts must not be orthogonal to existing concepts but they must have a certain correspondence / association to existing ones. This enables us to understand new concepts in terms of established concepts.

3.3.2 Conceptual anchoring non-spatial concepts

If we assume that the early concepts we form about the world are spatial concepts rooted in the physical environment through perception, it makes sense that we form subsequent concepts on the basis of spatial concepts, in particular concepts that are experientially related to spatial concepts – for example, temporal concepts. In our perceptual experience, there is a strong correlation between duration and spatial distance: longer distances usually take more time to travel than shorter distances, and we use expressions like ‘it takes long to get there’ which make sense both in a spatial and in a temporal interpretation.

But time also passes without any distance being traveled. The differences between physical distances we can measure with a yard stick and the experience that time can pass in which we can play, think, dream, or wait without physically moving make it worthwhile to form a new concept ‘duration’. Nevertheless, it is very hard, if not impossible, to describe temporal duration if we cannot resort to well-understood concepts like spatial distance that can be objectively characterized with reference to the physical environment. For example, we can suggest that a given duration is the time it would take to travel a certain distance by certain means; then we can abstract from the traveling event; everything else being equal, the duration persists. In fact, we may imagine temporal durations in terms of distances traveled to get a better grasp at the concept of duration.

3.3.3 The constructive approach in cognition

Other concepts like ‘speed’ may be developed accordingly (which concept is more basic, ‘duration’ or ‘speed’?) and when we have a good handle at everyday physical concepts, even more abstract concepts may be formed with reference to the existing ones. In their theory of metaphors, Lakoff and Johnson [15] give plenty of examples that illustrate relationships between different concepts; such examples may serve to build up abstract concept structures from concrete perceptual entities and relations. A great feature of such conceptual structures is that we may be able to borrow calculation or reasoning procedures from well-established concepts for use in new conceptual structures when we lack appropriate specific inference rules for the new concepts. For example, we might add durations as we add spatial distances without formally establishing the adequacy of this operation in the temporal domain, because this may be experientially plausible on the basis of our experience or knowledge about the spatial domain. In this way, conceptual systems including their inferential mechanisms can be bootstrapped from basic principles.

4 SPATIAL COGNITION IN AI

After providing arguments why spatial concepts are uniquely fundamental for cognition, I will outline how spatial concepts have been formalized for modeling and understanding spatial cognition in AI research.

4.1 Basic concepts

A fundamental insight into concept formation in cognitive systems is that new concepts are formed if distinctions are helpful or necessary. Conversely, the distinction of concepts or notions which do not make a significant difference in the given context is avoided. This appears to be an element of cognitive economy. For

knowledge representation in AI this means that concepts should be represented at the coarsest feasible rather than on the finest possible level.

The principle of using coarse representations appears to be in conflict with an engineering principle demanding the use of the most precise information available. In cognitive systems this conflict surfaces when we face the question how to represent perceptual information for spatial reasoning. After all, our eyes and other visual sensors provide rather precise information about the spatial environment over the entire perceptual range. In computer vision, the engineering approach has been followed for a long time: ‘why should we throw away information our sensors provide’ was a common statement in computer vision. Current developments with almost arbitrary availability of storage capacity tend to support this view.

Nevertheless, I propose a different approach. Our research question is: ‘which information is necessary to answer a given question?’ This question is geared towards unveiling cognitive principles – including the limits of their operation. In addition, we investigate how to exploit redundant knowledge to make cognitive systems robust.

4.2 Qualitative reasoning

Qualitative reasoning has been a topic in AI for a quarter century. ‘Qualitative’ is contrasted here to ‘quantitative’ or ‘metric’. Qualitative representations distinguish conceptual categories rather than measures. The simplest qualitative classification is obtained by comparing measurements in terms of *less than* (<), *equal* (=), and *greater than* (>). The basic idea of qualitative reasoning is that we use low-resolution representations to describe the essence of the state of affairs. From a cognitive perspective, this difference has important implications: while quantitative approaches need not care about the relation of the values to the represented environment, qualitative values depend on their meaning and their role in the represented world.

It may be interesting to note that in the empirical sciences, qualitative descriptions are considered inferior to quantitative descriptions; this is, of course, due to the fact that it is much harder to develop models that make quantitative predictions than models that provide qualitative descriptions; if I have quantitative information I can derive the qualitative information.

In qualitative reasoning, the situation is different: as we need representations not only to describe values but also to infer new configurations, we cannot rely on the translation of arbitrary numerical values into qualitative categories; instead, the main focus is on developing suitable conceptual frameworks that allow for inferring meaningful information from meaningful premises.

This suggests that we might need a new representation for each new problem area, as for each kind of task the relevant qualities may differ. To overcome this problem, we must identify general principles for qualitative reasoning that can be useful for a wide scope of problems. Where can we find such general principles? Right! – If we believe that space and time provide general structures for cognition as these are needed for interfacing percepts and concepts, we should expect to find these structures in space and time – if anywhere. It therefore may not be surprising that space and time were among the first domains successfully explored for qualitative reasoning.

4.3 Qualitative temporal reasoning

In our commonsense conceptualizations, the temporal domain appears simpler than the spatial domain: the simplest conception is that of a directed time-line with totally ordered time points. Time points correspond to temporal instants. Usually, we are more interested in events than in instants and events stretch over a time span. The philosopher Hamblin became interested in temporal relations between intervals. He described thirteen mutually exclusive relations between intervals that exhaustively characterize the relations that can hold between two temporal intervals [14]. Ten years later, Allen independently discovered these relations and developed a calculus of qualitative temporal relations [1]; the thirteen relations became widely known as ‘Allen relations’ in the AI community. Allen’s calculus has become a wonderful tool to reason about temporal relations of events on arbitrary levels of event hierarchies and across levels – and for this, only comparisons of meaningful entities (events) and no quantitative (external) measurements are required [9].

4.4 Qualitative spatial reasoning

If space is cognitively more fundamental than time, we need calculi for spatial reasoning. The idea is to process spatial information directly without deriving spatial notions from more general, abstract notions. Geometry allows us to compute spatial relations from quantitative measurements. In qualitative spatial reasoning we want to infer spatial relations on the basis of qualitative concepts instead.

Like temporal events, spatial entities form hierarchies: neighboring spatial objects can be aggregated to form larger spatial objects or spaces. Therefore, the spatial domain appears to be suitable for the exploitation of general spatial conceptualization and computing principles. As it turns out, there are many more meaningful ways to conceptualize and qualitatively represent commonsense space than time: space can be conceptualized in terms of three spatial dimensions, in terms of topological and ordering relationships, in terms of distance and orientation relations. In addition, we employ a variety of reference systems: absolute, intrinsic, relative, to use Levinson’s classification. [16].

Thus, we face the challenge of dealing not with THE representation of space but with many representations of space and their interrelationships. We must investigate the integration of spatial representations and processing mechanisms and the use of spatial knowledge as a resource for solving complex spatial problems [3].

5 SPATIAL AND CONCEPTUAL NEIGHBORHOOD

Neighborhood is a very important notion that helps us to structure and to comprehend the world around us and to help us to understand what is going on in it. Neighborhoods come into existence by grouping nearby objects and possibly separating them from other objects. On the discrete level of concepts, neighborhood corresponds to continuity on the geometric or physical level of description: continuous processes or events map onto identical or neighboring classes of descriptions.

Neighborhoods are interesting information structures, as neighborhoods can be formed recursively and represented by hierarchical tree or lattice structures.

In its most basic form, a hierarchy is a spatial structure, and the constituting nearness is a spatial distance. Spatial neighborhoods are very natural perceptual and cognitive entities as they are formed ‘automatically’ in front of our eyes when we move away from spatial configurations of objects. This is, because perceptual systems have limited spatial resolution and merge spatially neighboring objects. This is supported by the fact that perception organs themselves are organized spatially – neighboring retinal cells are sensitive to spatially neighboring stimuli. This spatial organization is maintained and supported by many neural structures in brains.

Neighborhoods also play an important role beyond the purely spatial domain. Other neighborhood structures can be derived from spatial neighborhoods, for example temporal neighborhood. In the context of reasoning, ‘conceptual neighborhood’ is particularly interesting. Conceptual neighborhood is defined as a structure on relations (rather than on objects or their locations): two (spatial, temporal, or other) relations are conceptual neighbors if they can be directly transformed into one another by a minimal change in the underlying world (i.e., if no other intermediate relation will hold in between).

Conceptual neighborhoods reflect operations and transitions in the represented reference world: only transitions that are possible in the reference world become conceptual neighbors. Thus, a world of rigid objects has a different conceptual neighborhood structure than a world with elastic objects or a world in which objects can be beamed to remote locations without reaching intermediate locations. In other words, the conceptual neighborhood structure of an environment maintains the (physical or geometric) laws that govern this environment. In this way, it maintains the integrity of the represented world and therefore it is a good basis for reasoning about operations and events in this world.

A main feature of conceptual neighborhood structures is that they can be used to control the complexity of reasoning processes: if the reasoning follows the structures established by the operations in the world, expensive computational operations can be avoided that would be carried out in a more abstract reasoning system. In some cases, a lower complexity class may be achieved through neighborhood reasoning [8, 9, 18].

6 SPATIAL AND NON-SPATIAL REPRESENTATIONS OF SPACE

In this paper I have argued in favor of representations and reasoning processes for cognitive systems that maintain and make use of spatial structures. So far, I discussed internal representations of spatial environments; I will now turn to external representations that are equally important for cognitive systems.

6.1 External representations of spatial knowledge

Intelligent processes are not only due to intelligent reasoning processes but also the result of successful interaction of cognitive systems with their environment. In particular, cognitive agents can resort to ‘knowledge in the world’ when they have smart sensory systems that can access relevant information on demand. Thus, the environment itself may partly replace representations that would be otherwise required [6]. The spatial environment itself features all the nice spatial properties we discuss in this paper.

Using the environment instead of a representation has advantages and disadvantages. The greatest advantage of the ‘representation in the world’ is that it is the most truthful and complete knowledge source we can possibly have; it is ‘merely’ a matter of perception to access this knowledge. Another advantage is that the knowledge is intrinsically related to the objects concerned.

On the other hand, ‘knowledge in the world’ is unprocessed information that does not contain any insights regarding the environment; it therefore is not knowledge in a strict sense. The intrinsic connection between objects in the world and information about these objects also can be a serious disadvantage: the knowledge may not be where the agent is and it may be hard for the agent to access the knowledge, as the environment may be large or inaccessible. Specifically, this concerns situations where the effort of planning is worthwhile, as abstract planning would be less expensive than trial and error in the real environment.

But there are intermediate avenues between ‘knowledge in the head’ and ‘knowledge in the world’ that can be taken. One option is to place preprocessed information in the world, for example symbolic classifications like street classifications or street names. Another option is the creation of a separate externalization of knowledge about the world that utilizes some of the advantages while avoiding some of the disadvantages of the ‘knowledge in the world’.

A well-known example of an external representation of a spatial environment is a geographic map: it maintains certain aspects of the spatial structure of the real environment while at the same time it avoids some of the disadvantages: map information should be more easily accessible than the real environment; due to scaling it may provide a better overview of the environment than the environment itself; it may contain preprocessed knowledge – in particular: aspectualized knowledge that emphasizes certain aspects and ignores others [4].

6.2 Linguistic representations of spatial knowledge

While spatially structured knowledge has numerous advantages, we must not forget the advantages of more abstract propositional (e.g. linguistic) representations. While spatially structured representations can express some uncertainty, they admit little ambiguity. Language, on the other hand, is capable of expressing ambiguity, conflicting information, and meta-knowledge rather easily. The capability of overcoming spatial constraints is a feature that may be particularly desirable, in certain situations.

Intrinsically spatial representations and non-spatial representations complement each other very well. Good examples for this are geographic maps: imagine a map without labels or a legend (they are suitable only to test spatial knowledge, but not to provide spatial knowledge) or legends / city street indices without maps (these are useful only to find out about the existence of streets in a city, but not about their spatial location).

A question sometimes asked about wayfinding is “what is better, a map or a linguistic route description?” This question can be answered rather well in the context of particular circumstances. For example, maps have much more general use as they represent a large number of routes while linguistic descriptions usually represent one particular route; on the other hand: maps need to be interpreted while linguistic descriptions are suitable to represent interpretations such that instructions can be directly executed [10].

Nevertheless, the question for the better of two alternatives erroneously suggests that I can either have one or the other, but not both. However, the best representations frequently are integrated hybrid representations that combine the advantages of various approaches. For example, a map provides the spatial background with a strong structural correspondence to the target domain while the linguistic information may provide useful meta-information that connects to the spatial background but uses its own instruments to be most effective.

7 SPATIAL SCHEMATIZATION

In this last chapter, I will discuss a form of spatial representation that combines features of intrinsically spatial forms of representation with features of linguistic representations.

Schematic maps like subway maps maintain certain aspects of spatial structures while relaxing others. While all geographic maps distort spatial relations of the environment to some extent due to representational constraints in the spatial medium map, schematic maps *intentionally* distort spatial relations to emphasize or deemphasize certain features of the represented environment. Specifically, topological relations are fully maintained while distance and orientation information is relaxed. The main objective in the schematization process is the elimination of decision-irrelevant features, especially curves, to reduce the maps to the essential structures that will be helpful for solving specific tasks – in particular wayfinding tasks. The goal is to reduce the cognitive load on the map user by eliminating distracting features.

Schematic maps are well suited to represent qualitative spatial concepts. The orientation of a line on the map may correspond to a general orientation (or category of orientations) in the environment; a distance on the map may correspond to the number of train stops, rather than to the metric distance in the environment, etc. [4, 13].

If we consider a physical spatial environment as constituting one extreme of a hypothetical continuum and abstract mental concepts of the spatial world as the other extreme, we can sequence different types of representations of spatial domains. Moving from the physical spatial environment towards abstract concepts, we obtain a mild abstraction by taking a visual image (e.g. a photograph) that preserves important spatial relations. Moving a few steps further, we may get a topographic map in which objects have been identified and spatial relations from the real environment are maintained. Further abstraction may lead to a schematic map.

Moving from the other extreme, abstract mental concepts can be manifested most easily by verbal descriptions. When we move in the hypothetical continuum closer to the physical manifestation of the world, we can put concepts of spatial objects and relations into a sketch map to convey selected spatial relations. Sketch maps tend to have close correspondences to verbal descriptions and they are used to augment verbal descriptions by spatial configurations from the physical world.

In this framework, schematic maps differ from sketch maps in that they are derived from topographic maps that are meant to represent a certain part of the environment completely, at a given granularity level. Sketch maps, on the other hand, usually correspond to the linear flow of speaking and drawing and frequently to the temporal sequence of route traversal. Thus, schematic maps provide information about a region while sketch maps more typically provide information about a single route or

about a small set of routes. However, there is no sharp boundary between schematic maps and sketch maps as schematic maps may be incomplete and sketch maps may be unusually elaborate.

8 CONCLUSION

Spatial cognition is a fundamental cognitive capability which serves as conceptual foundation for other cognitive capabilities. Through the connection between a variety of specific spatial sensors on one hand and general processing structures on the other hand, spatial cognition establishes an interface between physical spatial environments and abstract information structures.

Spatial cognition concerns the entire sensory – cognition – motor – environment loop including the registration of spatial relationships, the internal and external representation of spatial knowledge, the abstraction of spatial knowledge, the classification and matching of spatial and abstract structures, the transformation between spatial reference systems, the adaptation to effectors that act in and manipulate the spatial environment, and the interaction between cognitive agents. Thus, it can be viewed as exemplary for general cognition processes. The fact that spatial knowledge structures have strong links to physical structures that can be accessed and evaluated by objective means supports the use of scientific research methods.

The spatial domain also is particularly interesting as a domain for exploring imprecise, incomplete, uncertain, coarse, and conflicting knowledge. These attributes can be related to physical perception processes; this allows us to provide physical explanation models for the formalization of this kind of knowledge. The physical domain can play a similar role in cognitive systems as logics plays in formal systems, as our understanding of simple physical systems is comparable to our understanding of simple logic systems.

As natural cognitive systems are understood to a degree where at least partial functionality can be replicated, powerful computer systems are available that can keep up with natural computation, on certain levels, and robot and multimodal interaction devices become more and more sophisticated, high-level spatial cognitive functions can be explored by a constructive approach and new capabilities can be integrated into artificial cognitive agents.

If we maintain a close relation between human spatial concepts and spatial concepts we use for our artificial agents, we support the communication between human and artificial cognitive agents. In this way, we support the development of spatial assistance systems that complement rather than replace human cognitive abilities.

ACKNOWLEDGMENTS

I acknowledge critical and constructive comments by Thomas Barkowsky. I am grateful for support for the Transregional Collaborative Research Center *Spatial Cognition: Reasoning, Action, Interaction* (SFB/TR 8) provided by the Deutsche Forschungsgemeinschaft. I am also grateful for support for the International Quality Network on *Spatial Cognition* provided by the German Academic Exchange Service (DAAD).

REFERENCES

[1] J.F. Allen, Maintaining knowledge about temporal intervals, *CACM* 26 (11) (1983) 832-843.

- [2] J. Ayers, J.L. Davis, and A. Rudolph, eds., *Neurotechnology for biomimetic robots*, MIT Press, Cambridge, MA, 2002.
- [3] T. Barkowsky, The Transregional Collaborative Research Center, *KI* 1/2004, 34-36.
- [4] B. Berendt, T. Barkowsky, C. Freksa, and S. Kelter, Spatial representation with aspect maps, in C. Freksa, C. Habel, and K.F. Wender, eds., *Spatial cognition – an interdisciplinary approach to representing and processing spatial knowledge*, 313-336, LNAI 1404, Springer, Berlin 1998.
- [5] V. Braitenberg, *Vehicles: Experiments in synthetic psychology*, MIT Press, Cambridge, MA, 1984.
- [6] R.A. Brooks, Intelligence without representation, *Artificial Intelligence* 47, 1991, 139-159.
- [7] C. Freksa, Qualitative spatial reasoning, in D.M. Mark and A.U. Frank, eds., *Cognitive and linguistic aspects of geographic space*, 361-372, Kluwer, Dordrecht 1991.
- [8] C. Freksa, Conceptual neighborhood and its role in temporal and spatial reasoning, in M. Singh, L. Travé-Massuyès, eds, *Decision support systems and qualitative reasoning*, 181-187, North-Holland, Amsterdam 1991.
- [9] C. Freksa, Temporal reasoning based on semi-intervals, *Artificial Intelligence* 54 (1992) 199-227.
- [10] C. Freksa, Spatial aspects of task-specific wayfinding maps: A representation-theoretic perspective, in J.S. Gero and B. Tversky, eds., *Visual and Spatial Reasoning in Design*, 15-32, Key Centre of Design Computing and Cognition, University of Sydney, 1999.
- [11] C. Freksa and T. Barkowsky, On the relation between spatial concepts and geographic objects, in P. Burrough and A. Frank, eds., *Geographic objects with indeterminate boundaries*, 109-121, Taylor and Francis, London 1996.
- [12] C. Freksa and C. Habel, Warum interessiert sich die Kognitionsforschung für die Darstellung räumlichen Wissens? in C. Freksa and C. Habel, Hg, *Repräsentation und Verarbeitung räumlichen Wissens*, 1-15, Springer, Berlin 1990.
- [13] C. Freksa, R. Moratz, and T. Barkowsky, Robot navigation with schematic maps, *Intelligent Autonomous Systems* 6, 809-816, IOS Press, 2000.
- [14] C.L. Hamblin, Instants and intervals, in: J.T. Fraser, F.C. Haber, and G.H. Müller, eds., *The study of time*, Springer, Berlin, 1972, 324-331.
- [15] G. Lakoff and M. Johnson, *Metaphors we live by*. University of Chicago Press, Chicago, 1980.
- [16] S.C. Levinson, Frames of reference and Molyneux's question: crosslinguistic evidence, in P. Bloom, M. Peterson, L. Nadel, and M. Garrett, eds, *Language and Space*, 109-169. MIT Press, Cambridge, MA 1996.
- [17] D.M. Mark, C. Freksa, S.C. Hirtle, R. Lloyd, and B. Tversky, Cognitive Models of Geographic Space, *Int. J. of Geographical Information Science* 13 (1999), 8, 747-774.
- [18] B. Nebel and H.-J. Bürckert: Reasoning about temporal relations: a maximal tractable subclass of Allen's interval algebra. Proc. AAAI 1994: 356-361.
- [19] R. Pfeifer and C. Scheier, *Understanding intelligence*, MIT Press, Cambridge, MA, 1999.
- [20] D.E. Rumelhart and D.A. Norman, Representation in memory. In R.C. Atkinson, R.J. Hemstein, G. Lindsey, and R.D. Luce, eds., *Steven's handbook of experimental psychology*, 511-587, 2nd ed., John Wileys & Sons Ltd, New York, 1988.