Spatial Computing

How Spatial Structures Replace Computational Effort

Christian Freksa

Abstract At the Advanced Study Institute on Cognitive and Linguistic Aspects of Geographic Space in Las Navas del Marqués in July 1990, I presented a chapter on *Qualitative Spatial Reasoning*. In that chapter, I suggested that spatial inference engines might provide the basis for rather general cognitive capabilities inside and outside the spatial domain. In the present chapter, I will follow up on this perspective and I will illustrate the ways in which research in spatial cognition has progressed towards understanding spatial reasoning and spatial computing in a more literal sense: using a spatial substrate. The chapter presents a progression of approaches to spatial reasoning from purely descriptive to increasingly spatially structured. It demonstrates how spatial structures are capable of replacing computational processes. It discusses how these approaches could be developed and implemented in a way that may help us to better understand higher-level spatial abilities of cognitive systems that are frequently attributed to the right cerebral hemisphere in humans. The chapter concludes by discussing the special role of space and time for cognition and advocates a thorough overall analysis of the specific problem to be solved to identify the most suitable approach to computation.

Keywords Spatial cognition · Spatial computing · Symbolic computing · Structure of space · High-level spatial abilities

C. Freksa (✉)
Spatial Cognition Research Center SFB/TR 8, University of Bremen, Bremen, Germany
e-mail: freksa@sfbtr8.uni-bremen.de

M. Raubal et al. (eds.), *Cognitive and Linguistic Aspects of Geographic Space*,
Lecture Notes in Geoinformation and Cartography, DOI: 10.1007/978-3-642-34359-9_2,

1 Spatial Problems

Let us consider examples of common spatial problems we might encounter. The problems are to be taken literally; i.e., no additional information is provided; missing information must be added from knowledge or assumptions about the environment.

1. Given the triangle *ABC* with the coordinates $A = (1, 3)$, $B = (9, 2)$, $C = (6, 8)$; is $P = (8, 4)$ inside or outside the triangle *ABC*?
2. (How) can I get the piano into my living room?
3. How do I get from here to John's place?
4. Which is closer: from here to John, or to Mary?
5. Is the tree on my property or on your property?
6. 

Problem 1 is a classic high school geometry problem which can be solved abstractly with linear equations; the correct algebraic solution will locate *P* on the line *BC*; numeric solutions may place *P* inside or outside the triangle, depending on the number format and algorithm chosen.

Problem 2 is a form of the classic *Piano Movers' Problem* in mathematics (Schwartz and Sharir 1983); although this problem can be represented geometrically, in practice it is rarely approached mathematically in the abstract representation domain but by trial and error in the physical problem domain.

Problem 3 cannot very well be presented in geometric terms; a graph structure that depicts the location 'here', John's place, and a traversable connection between them is more appropriate and often times preferable to a solution in the physical domain, particularly if John's place is far away.

Problem 4 typically does not require the mathematically correct solution, which may take a long time to determine. A quickly provided estimate tends to be more helpful, in practice.

Problem 5 is an example where a formal approach alone may not suffice. Although the boundaries of the properties will be defined in a legal document in terms of precise geo-coordinates, the real-world correspondence of the legal boundary may be too expensive to determine; therefore the correct boundary often is not known. In addition, it may not be clear where the branches and roots of the tree start and where they end, in formal terms.

Problem 6 (related to the Piano Movers' Problem) is not posed in terms of words or numbers but in terms of spatial objects (resp. an image thereof). It is a truly spatial problem presented physically to small children who will try to fit the small colored objects into the openings of the wooden cube and thus learn about spatial features like size and shape through physical processes by trial and error.

These examples illustrate that spatial problems may come in different modalities: in terms of numbers, language, or spatial configurations; and in different domains: abstract mathematical or legal and concrete physical space. Likewise, the solutions to spatial problems may be required in terms of numbers, language, or spatial configurations. The solution may or may not be needed in the same modality or domain as the problem statement. A correct solution may not always be the best solution, as quickly or cheaply available sub-optimal solutions may be more useful in certain situations. In other words, we may need to transform problems and solutions between different modalities and domains, and the generation of a problem solution may take place in a variety of modalities and domains. Accordingly, it may be helpful to have approaches available that are tailored to the respective requirements (cf. Sloman 1985).

This observation raises the issue whether we always have to transform spatial problems into geometric formalisms to enable computational solutions by means of sequential interpretation of instructions, or whether we can find ways to directly process entire spatial configurations, as humans seem to be able to do (Shepard and Metzler 1971). I will dub the classic computer science approach of sequential interpretation as *left-brain computing,* as information processing in the left cerebral hemisphere is associated with bottom-up or language-like sequential processing; I will dub the approach of processing entire spatial configuration as *right-brain computing,* as the right cerebral hemisphere in humans is associated with top-down or holistic processing (cf. Kosslyn 1987).

In the present chapter, I will first review foundations of qualitative temporal and spatial reasoning. I will then discuss the notion of *conceptual neighborhood* and how we can exploit this notion for spatial computing. I will introduce tools for processing qualitative spatial relations, and then address the transition from spatial relations to spatial configurations. Finally I will demonstrate and analyze the notion of *spatial computing* as contrasted to *symbolic computing*.

2 Qualitative Temporal and Spatial Reasoning

The starting point for much of the research in qualitative temporal and spatial relations since the late 1980s was the chapter *Maintaining knowledge about temporal intervals* by Allen (1983), although the underlying insights had been published previously (Nicod 1924; Hamblin 1972).

The intriguing result of this research was that thirteen 'qualitative' relations could describe temporal relations between events uniquely and jointly exhaustively (Fig. 1). There was an expectation that the idea of qualitative relations could

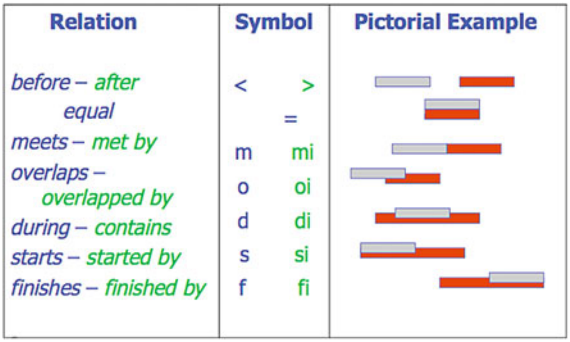

Fig. 1 The 13 jointly exhaustive and mutually exclusive qualitative relations between two temporal intervals

B r2 C / A r1 B	<	>	d	di	o	oi	m	mi	s	si	f	fi
"before" <	<	no info	< o m d s	<	<	< o m d s	<	< o m d s	<	<	< o m d s	<
"after" >	no info	>	> oi mi d f	>	> oi mi d f	>	> oi mi d f	>	> oi mi d f	>	>	>
"during" d	<	>	d	no info	< o m d s	> oi mi d f	<	>	d	> oi mi d f	d	< o m d s
"contains" di	< o m di fi	> oi di mi si	o oi dur con =	di	o di fi	oi di si	o di fi	oi di si	di fi o	di	di si oi	di
"overlaps" o	<	> oi di mi si	o d s	< o m di fi	< o m	o oi dur con =	<	oi di si	o	di fi o	d s o	< o m

Fig. 2 Upper part (facsimile) of the composition table for the qualitative temporal relations (without the 'equals' relation) from Allen (1983). The relation r1 is composed with the relation r2 to obtain the composite relation found in the table. In most cases, more than one relation may result from a composition. "no info" means that all 13 relations may result from a given composition

be extended to one- and higher-dimensional spatial objects that share the extendedness property of temporal intervals. Initially, researchers had in mind a single spatial calculus that would compute all-embracing spatial relations between objects based on information about spatial relations between other objects. However, it became apparent soon that it would be more effective to develop specialized calculi that deal with individual aspects of space rather than a comprehensive spatial calculus that would integrate multiple aspects of space in a single formalism. For example, Allen's interval calculus (see Fig. 2 for the rules of combining interval relations) can be easily adapted to 1-dimensional oriented space (Freksa 1991b; Skiadopoulos and Koubarakis 2004; Liu and Li 2011) or to three spatial dimensions individually (Guesgen 1989).

3 Conceptual Neighborhood

Interval relations can be described at a finer level of resolution in terms of point relations, i.e., in terms of relations between the starting points and the ending points of the intervals. An important feature of physical time and space is that gradual change in position or size results in small qualitative changes or no changes at all between the point relations involved. For example, in the transition from the *before* relation to the *meets* relation, only one of the four point relations between beginnings and endings of the intervals changes: the relation between the ending of the first interval and the beginning of the second interval changes from *smaller than* to *equals*. Accordingly, spatio-temporal configurations that result from small physical changes are perceptually and cognitively closely related.

Furthermore, events in close temporal vicinity are related more easily to one another than events in different epochs. Similarly, nearby spatial locations are more easily related to one another than locations far apart. This insight is captured in *Tobler's First Law of Geography:* "Everything is related to everything else, but near things are more related than distant things" (Tobler 1970).

The role of nearness extends from temporal and spatial neighborhood to the more abstract level of relations: certain relations are closer to one another than others; in fact, some relations are distinguished only by a single detail. These relations are called *conceptual neighbors* (Freksa 1991a). In Fig. 3 the thirteen interval relations from Fig. 1 are applied to one-dimensional oriented space. Conceptually neighboring relations are depicted next to each other.

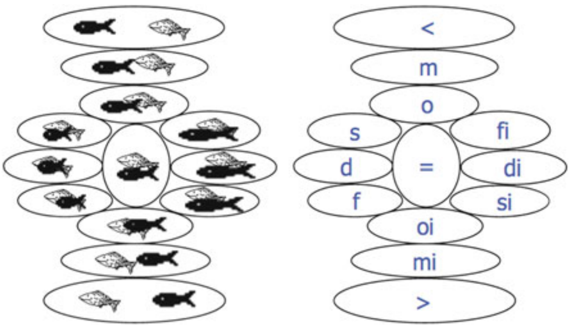

Fig. 3 *Left* Thirteen qualitative relations for objects (here fishes) in one-dimensional oriented space. The example classifies positions of objects in the horizontal dimension. The 13 relations are arranged by *conceptual neighborhood*. *Right* The corresponding labels of the qualitative temporal relations from Allen (1983) depicted in the same spatial arrangement (adapted from Freksa 1991b)

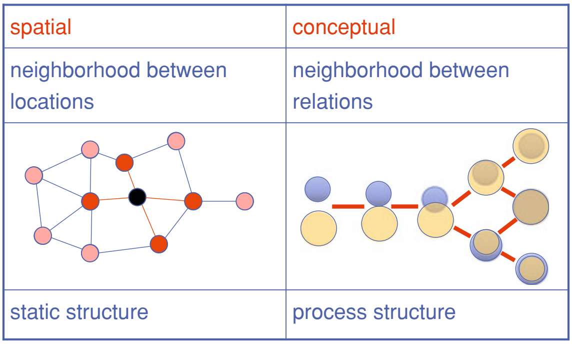

Fig. 4 Spatial and conceptual neighborhood: The *left* graph depicts relations between static spatial locations; directly connected nodes represent spatial neighbors. The *right* graph depicts direct transitions between spatial relations due to processes in the domain; edges correspond to conceptually neighboring relations caused by a minimal spatial change in the domain

The notions of conceptual and spatial neighborhood are closely related: Whereas two directly connected objects are called *spatial neighbors*, two by minimal differences directly connected relations are called *conceptual neighbors* (Fig. 4).

Arranging temporal and spatial relations by conceptual neighborhood enables numerous features for representing spatial knowledge and for spatial reasoning:

- Sets of neighboring relations can be lumped together to define *coarse relations* (Freksa 1992a, b);
- Conceptual neighborhoods define hierarchies for representing incomplete knowledge (Freksa and Barkowsky 1996);
- Qualitative reasoning based on conceptual neighborhoods allows for efficient non-disjunctive reasoning (Nebel and Bürckert 1995; Balbiani et al. 2000);
- Neighborhood-based incomplete knowledge can be easily augmented as additional knowledge is gained during successive reasoning (Freksa 1992b);
- Coarse relations based on conceptual neighborhoods frequently exhibit a natural correspondence to everyday human concepts (Freksa 1992a);
- Spatial and temporal inferences in qualitative reasoning typically result in conclusions that form conceptual neighborhoods (Freksa 1992a, b);
- Conceptual neighborhoods can be formed on various levels of granularity (cf. Fig. 5).

4 Neighborhood-Based Reasoning

One important feature of conceptual neighborhood-based abstraction is that *incomplete knowledge* can be conceptualized and represented as *coarse knowledge* (Fig. 5). By abstracting from missing or unnecessary details, reasoning can be carried

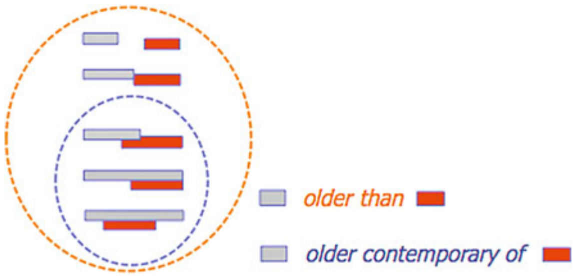

Fig. 5 Coarse temporal relations forming an abstraction hierarchy. The relation 'older contemporary of' corresponds to the conceptual neighborhood of the three finer relations 'overlaps', 'finished by', and 'contains'. The even coarser relation 'older than' corresponds to a larger conceptual neighborhood that additionally includes the two fine relations 'before' and 'meets'

out efficiently. In this way, computationally and conceptually problematic properties of disjunctive knowledge processing are avoided which are encountered when incomplete knowledge is represented as a set of completed potential alternatives.

Coarse reasoning does not necessarily yield coarser results than reasoning with fine relations. But reasoning with coarse relations calls for different inference procedures than reasoning with fine relations. Conjunctions of partially overlapping coarse inferences based on imprecise or incomplete knowledge fragments from different sources result in more precise or *fine* conclusions if the premises are appropriately chosen. With this property, the coarse reasoning approach is suited to model the synergy of multimodal coarse knowledge sources that result in precise knowledge (cf. distributed representations and *coarse coding* in biological or artificial perceptual systems, e.g. Edelman and Intrator 2000). Figure 6 presents a coarsened version of the Allen composition rules that exploits conceptual neighborhood relations between fine relations. For example, the relation *older contemporary* (oc) corresponds to the union of *overlaps* (o), *finished by* (fi), and *contains* (di).

5 A Multitude of Specialized Calculi and SparQ

A considerable variety of spatial calculi have been developed over the past 20 years (Cohn and Hazarika 2001); these can be classified as

- Measurement calculi, e.g. Δ-Calculus (Zimmermann 1995);
- Topological calculi, e.g. 4-intersection calculus, 9-intersection calculus, RCC-5, RCC-8 (Egenhofer and Franzosa 1991; Randell et al. 1992);
- Orientation calculi, e.g. point/line-based: DCC, FlipFlop, QTC, dipole or extended objects (Freksa 1992b; Ligozat 1993; Van de Weghe et al. 2005; Moratz et al. 2000);
- Position calculi, e.g. Ternary point configuration calculus (TPCC—Moratz et al. 2003).

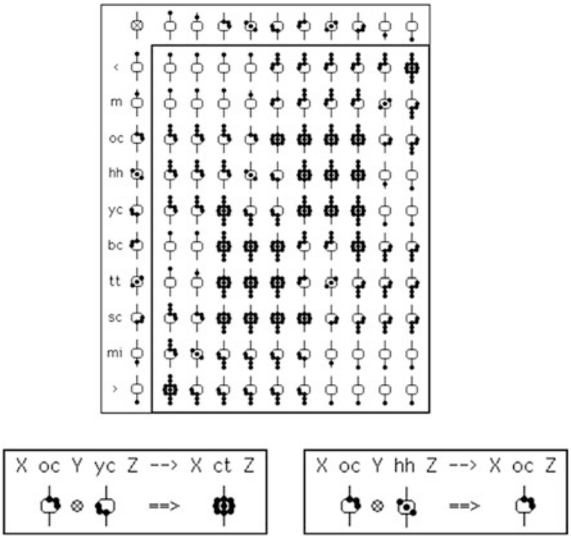

Fig. 6 Conceptual neighborhood-based composition table and inferences based on this table. Spatial pictograms symbolically depict 1D oriented spatial or temporal relations. Each black dot corresponds to a fine relation; conceptually neighboring relations form lumps of dots that correspond to coarse relations. *Top* Conceptually neighboring columns and conceptually neighboring rows from the original table have been merged. *Bottom* Two coarse inferences using this composition table (composition operator is denoted by $\otimes$). Above the pictograms, the relations are symbolized in classical logic notation. For an elaborate explanation see (Freksa 1992a)

To simplify and support the use of qualitative spatial calculi for specific reasoning tasks, various tools have been developed. Prime examples are SparQ (Wallgruen et al. 2007; Wolter and Wallgruen 2012); GQR (Westphal et al. 2009); QAT (Condotta et al. 2006); and CLP (QS) (Bhatt et al. 2011). While some approaches focus at specialized spatial reasoning methods, others aim to integrate specialized techniques with general knowledge representation methods for logic-based reasoning.

The toolbox SparQ[1] integrates numerous calculi for qualitative spatial reasoning and allows for adding arbitrary binary or ternary calculi through the specification of their base relations and their operations in list notation or through algebraic specification in metric space. SparQ has a modular architecture and can easily be extended by new modules (Fig. 7).

SparQ performs a number of operations that are helpful for dealing with spatial calculi:

- Qualify: quantitatively described configurations are translated into qualitative relations;

[1] www.sfbtr8.spatial-cognition.de/project/r3/sparq/ (accessed: 1 Jan 2012).

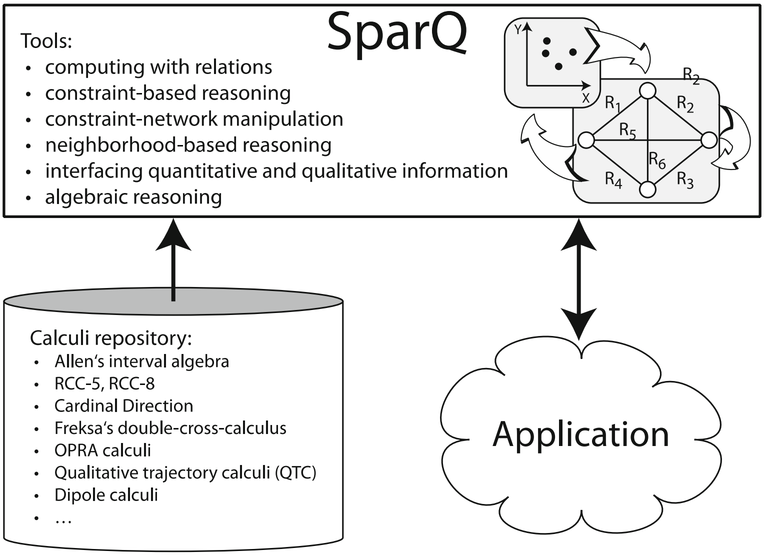

Fig. 7 Modular SparQ architecture. Operations in different qualitative reasoning calculi can be invoked through standardized commands (from Wolter and Wallgruen 2012)

- Compute-relation generates a qualitative inference for a given calculus based on the premise relations and the calculus specification;
- Constraint-reasoning allows for the specification of an inference strategy on a given spatial configuration and returns scenarios that are consistent with the configuration; if the description of the scenario is inconsistent, SparQ informs about the inconsistency;
- Neighborhood-reasoning enables conceptually compatible constraint relaxation and yields semantically meaningful neighboring inferences;
- Quantification generates prototypical 'general' pictorial instances of abstract qualitative descriptions (this is still in an experimental stage).

Although it is helpful to have a variety of calculi available in uniform specification and interface languages, there is still an issue about which calculus to select to solve a given problem. Thus, there is a challenge to understand and describe spatial calculi on the meta-level. The goal is to specify spatial configurations and the type of required problem solution in such a way that the available calculi can be automatically configured to solve the problem.

6 From Spatial Relations to Spatial Configurations

Quantitative computation of spatial configurations by means of Euclidean geometry is well understood. For example, in planar geometry, we can compute all angles, heights, and the area of arbitrary triangles, if the lengths of the edges of the triangles are given by means of the formulae depicted in Fig. 8.

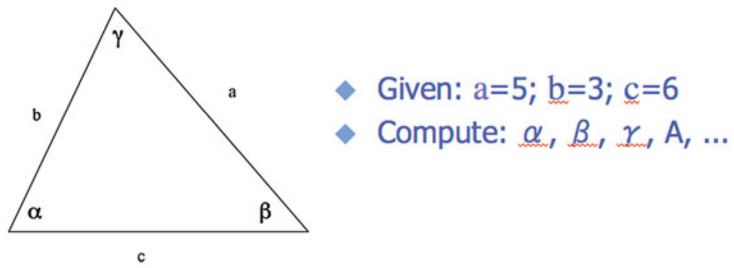

$$\alpha = \arccos\left(\frac{b^2+c^2-a^2}{2bc}\right)$$
$$\beta = \arccos\left(\frac{a^2+c^2-b^2}{2ac}\right)$$
$$\gamma = \arccos\left(\frac{a^2+b^2-c^2}{2ab}\right)$$

$$A = \frac{1}{2}ah_a = \frac{1}{2}bh_b = \frac{1}{2}ch_c,$$

$$h_a = c \cdot \sin\beta = b \cdot \sin\gamma$$
$$h_b = a \cdot \sin\gamma = c \cdot \sin\alpha$$
$$h_c = b \cdot \sin\alpha = a \cdot \sin\beta$$

Fig. 8 Formal abstraction of geometric relations in the Euclidean plane

These formulae are valid for planar spatial configurations independently of position, orientation, scale, or other influences. Spatial relations in physical environments conform to topological and geometric laws that are not affected by contextual influences from other modalities. As a consequence, only few constraints need to be specified and many—or even all—spatial relations are determined.

The principle is well known from high school geometry. For example, on a flat sheet of paper, we can construct exactly two triangles from the specification of three line segments, provided the specified lengths conform to the triangle inequality. In this construction, a compass and ruler are capable of qualitative representation and they exhibit certain abstraction capabilities: the compass represents a distance equal to the length of a given line segment and can apply this distance abstracting from location and orientation. Similarly, the ruler represents a distance and can apply it to any pair of points, independently of orientation and location (within practical bounds).

7 Preserving Spatio-Temporal Structure

Although the formal abstraction shown in Fig. 8 is capable of generating arbitrary spatial relations through abstract computation, the abstraction mechanism does not preserve spatial structure in the way neighborhood-based representations preserve the structure of the represented spatial domain. Structure-preserving representations exploit structural correspondences between the representation medium and the represented domain. They have the advantage that essentially the same operations can be applied to the representation as to the represented domain. For example, on a geographic map we can navigate much like in the geographic environment with the advantage that we can maintain an overview more easily and that we do not need to cover large distances.

As a consequence, structure-preserving representations (Sloman 1971) are advantageous at least for those situations in which humans use the representations; this is the case for assistance systems, for example, where spatial and temporal representations are employed as human–machine interfaces. Humans can carry out zooming operations by moving towards or away from the representation medium; at the same time they can perform refinement and coarsening operations; they can perform perspective transformations by looking at the medium from different angles; they can aggregate and partition spatial regions by making use of natural neighborhood structures; they can move across the medium much like in the represented domain and they can experience spatial and conceptual transitions while doing so; structure-preserving media also may support shape transformation operations in similar ways as in the represented domain.

Are there additional reasons for exploring structure-preserving representations besides the convenience for human users? I believe so. The operations described in the previous paragraph are helpful not only for human users; they may be useful whenever

- problem statement and problem solution are in the spatial domain;
- there is a single spatial configuration about which we may want to answer many questions;
- there are agents with spatial perception and locomotion, e.g., mobile robots;
- several agents need to communicate about a given spatial configuration;
- they can save resources by avoiding unnecessary operations.

In other words: structure-preserving representations also may be advantageous for machine processing. We will come back to this consideration in the next section.

Geometric-diagrammatic constructions on a piece of paper can serve as structure-preserving representations of space, since flat paper provides the universal spatial structure that guarantees the correctness of trigonometric relations in a planar domain. Figure 9 depicts universal correspondences between geometric functions in planar spatial structures.

Computation by diagrammatic construction is a form of *analogical reasoning* (cf. Gentner 1983): the basis for establishing analogies is given through the universal spatial interdependencies that justify the comparison between the source domain and the target domain; the analogies usually concern the abstraction from specific values in the domain. Nevertheless, geometric constructions are sequential constructions that are most easily described by classical algorithms and procedures.

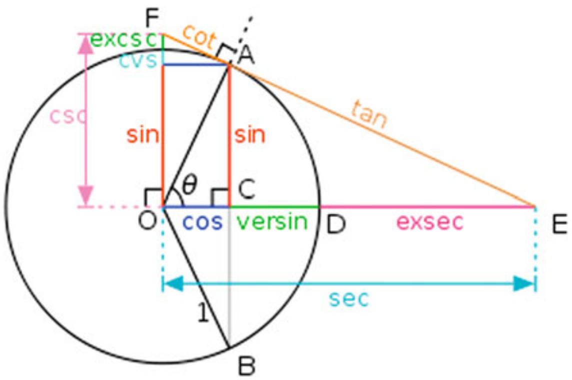

Fig. 9 Spatial construction of trigonometric functions. The graph depicts interdependencies of geometric relations. All trigonometric functions of an angle Θ can be constructed geometrically in terms of a unit circle centered at O

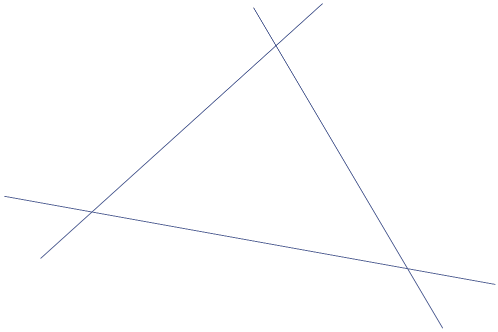

Fig. 10 Three line segments are applied to a spatially structured domain. Numerous new entities and relations are established through the interaction of these lines and the constraints of the domain: nine new line segments, 12 angles, a triangle, its area, and the spatial relations between all these entities

8 Space as Computer

In his book *Rechnender Raum* ('Computing Cosmos' or 'Calculating Space') (Zuse 1969), the computer pioneer Konrad Zuse discussed the issue of structure correspondence between computational representations and the physical domain. He addressed the issue on the micro-level of discrete versus continuous structures, maintaining that discrete representations only approximate continuous structures and mimic random deviations rather than replicating the physical laws of quantum mechanics.

In this section, I want to discuss the idea of structure correspondence on the macro-level of spatial configurations and carry the notion of diagrammatic construction one step further.

Suppose we apply three line segments to a flat surface as shown in Fig. 10. What do we see in this figure? We can easily identify nine additional line segments of specific lengths, three line intersections at specific locations, twelve specific pairwise identical angles, one triangle with a specific area, and numerous relations between those entities.

Where did all these entities and relations come from as we only placed three simple straight lines onto the surface? One way to answer this question is: The surface *computed* these entities and relations according to the laws of geometry. This would be the type of answer we would give if we gave a computer the line equations and the procedures to generate the mentioned entities and relations. What is the difference between the computer approach and the ‘flat paper approach’?

The computer algorithm encodes knowledge about the spatial structure of the surface that enables its interpreter to reconstruct in a sequential procedure step-by-step certain abstractions of its spatial structure that are constrained by abstract representations of the lines and their relationships. On the other hand, the flat surface itself and its spatial structure relate *directly* and *instantly* to the lines and generate the entities and relations without computational procedure by means of the inherent structural properties. It represents space rather than knowledge about space. This is why I call this approach *spatial computing* rather than knowledge processing.

9 The Notion of Spatial Computing

Much of what we do in artificial intelligence and computer science takes place on the knowledge level (Newell 1982). The hype of general purpose computing in the 1960s was based on the insight that we can express everything we can think and talk about in terms of physical symbols and that we can manipulate these symbols in a computer similarly to the way we reason and talk *about* arbitrary domains. In this way we can use our knowledge and understanding of these domains to answer questions about them and to solve problems. In the generality of this insight we may have lost sight of the fact that the domains of space and time are omnipresent not only in the worlds we talk about but also within our physical symbol manipulation systems. Considering this fact, couldn’t we make use of the spatial and temporal properties of these physical systems on the object level rather than reason about them on the meta-level? For certain tasks in the spatial and temporal domains we would simply act in space and time and *see* what happens rather than process knowledge about space and time and *know* what happens.

Figure 11 schematically depicts the relation between the meta-level of formal and computational reasoning and the object level of spatial configurations. The formal reasoning approach to computing spatial relations is shown in the upper part of the figure; the approach that applies spatial structures directly is shown in the lower part of the figure. In the classical formal reasoning approach, the task either must be given in formal terms or it must be formalized from object-level configurations before formal reasoning processes can be invoked. The formal result can be presented on a formal level or be transformed to an object-level configuration by instantiation.

The spatial computing paradigm takes place on the object level of spatial configurations. The task is directly presented as spatial configuration (e.g., the configuration of Fig. 10), or a spatial configuration is generated through instantiation from a formal specification. If the task is to answer questions about spatial properties and relations of the configuration, the result is available instantly due to

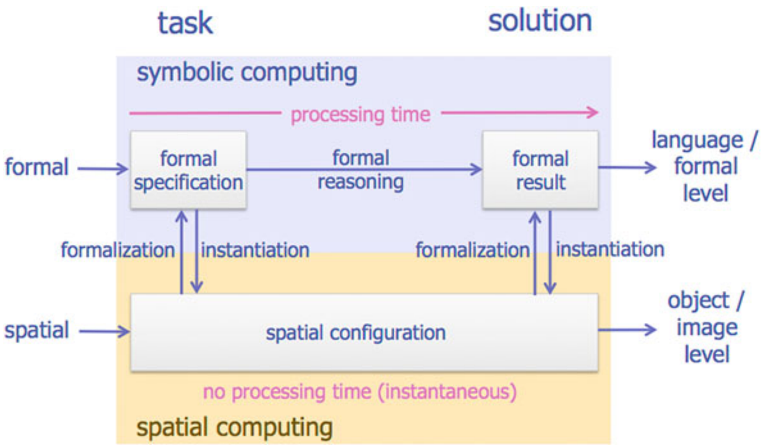

Fig. 11 Two approaches to generating spatial entities and relations: in the *upper* part of the figure, a classical sequential symbolic computing approach transforms a formal specification of a spatial configuration by means of formal reasoning into a formal result on the meta-level. In the *lower* part a spatially structured substrate on the object level guarantees compliance with spatial constraints and instantly makes available all spatial implications of the configuration in spatial form. Specific relations can be read-off directly from the configuration. Transformations between the object level and the meta-level can be carried out at the task stage or the solution stage

the intrinsic constraints of the spatial substrate on the object level. If the task involves physical operations on spatial configurations, these operations will be subject to spatial constraints of changing physical configurations and require processing time; but all spatial implications of the operations will be available instantly as the operations are performed. Depending on the form in which the result of the 'computation' is needed, we need processes that extract the desired result from the configuration (e.g., a line segment or an angle) as input for the next spatial computing task; if the result is needed on the formal level, the object-level entity needs to be formalized, e.g., for use in a classical computation process.

The spatial computing approach involves a paradigm shift that makes it difficult to compare with symbolic approaches by purely computational measures. The reason is that in the symbolic realm, we assume that problems are given in formalized form and that the results will be needed in formalized form, as well. Perceptual operations necessary to formalize spatial knowledge are not taken into account in symbolic approaches. Thus, computational cost is restricted to symbol manipulation processes. However, real-world spatial situations may be different as the example problems in the introduction suggest: for some of them the formalization task may be too time-consuming or expensive; a direct object-level mapping to a spatial substrate may become more feasible, particularly, as sensor technology continues to develop. Perception processes that had to be performed by humans in the process of formalizing spatial knowledge become a part of the spatial computing paradigm; however, they will not be discussed in this chapter.

The term ‘spatial computing’ is used by various researchers in interesting ways that are related to the topic of this chapter: In a Dagstuhl Seminar on *computing media and languages for space-oriented computation* the organizers state that ‘… it is important to make **space** not an issue to abstract away, but a first-order effect that we optimize. The distinguishing feature of *spatial computing* then is that computation is performed distributed in space and topology define the computation’ (DeHon et al. 2007). Similar goals are stated in a proposal for *spatial cloud computing* for use in the Geospatial Sciences (Yang et al. 2011). A student of the MIT School of Architecture and Planning posted the manuscript for a master thesis (Greenwold 2003) in which he defines ‘Spatial computing is human interaction with a machine in which the machine retains and manipulates referents to real objects and spaces.’ The artist Albert Hwang presents a video film series (Hwang 2012) in which he demonstrates simulations of augmented reality technology for interaction with spatial environments to illustrate the power of spatial computing technology that is yet to be developed.

Whereas the former two projects share technical goals—using spatial substrates for computation—with the approach presented here, the latter two projects share some of the motivation for our approach: interaction with spatial environments through perception and action. The main motivation for our approach, however, is to understand cognitive functionality by conceiving and implementing suitable representations, processes, and technical realizations.

10 Basic Entities of Cognitive Processing

In geometry, the spatial world can be described in terms of infinitesimally small points; lines are defined in terms of points; areas in terms of lines; etc. In contrast, in cognition, basic entities usually are not infinitesimally small points; they may be

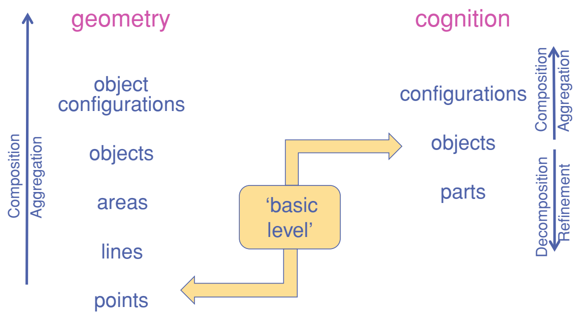

Fig. 12 Two ways to conceptualize physical objects. *Left* In geometry, we aggregate arbitrarily complex structures from atomic point entities. *Right* In cognition, basic entities may be geometrically complex, meaningful entities. Through cognitive effort, basic entities can be decomposed into more elementary entities or aggregated into more complex configurations; cognition on the level of basic cognitive entities is possible without invoking elementary constituents

entire physical objects like books or chairs (Fig. 12). Basic cognitive entities carry meaning related to their use and function and we perceive and conceptualize them in their entirety even if certain details are not accessible to our perception. We may not even know about the composition of basic cognitive entities; still, we are able to talk about them and to use them for daily activities. The cognitive apparatus appears to be flexible as to which level in a huge lattice of part-whole relations to select as 'basic level' (cf. Rosch 1978); it also appears to be able to focus either on the relation between an object and a configuration of objects, or alternatively, on the relation between an object and its parts. Both transitions involve cognitive effort, while the mere consideration of the basic level appears almost effortless.

It is known that we can apply simple mental operations, e.g., mental rotation, to entire spatial objects at once (Shepard and Metzler 1971). In spatial computing, we would like to implement processes that have comparable capabilities: manipulating entire objects without manipulating all their constituent parts. We would expect to obtain only a coarse result of cognitive operations on basic cognitive entities with little effort; to resolve details, we would have to invest cognitive power. This processing approach would be in contrast to geometric spatial processing where we would expect to know the details before we know the complex structure.

11 Conclusions and Outlook

Let me return to the spatial problems that I used in the beginning of the chapter to introduce various perspectives on spatial challenges. The main message of this exercise is that spatial problem solving consists of more than solving equations. First of all, a spatial problem needs to be perceived as one. Second, it needs to be represented as one. Third, the representation needs to be processed. Fourth, the result needs to be interpreted in spatial terms.

With regards to the representation of spatial knowledge for problem solving we have lots of options, as there are many ways to conceive of space. For example, space may be conceived of as empty space—"what is there when nothing is there"—or as the space spanned by physical objects. Space can be described in terms of a multitude of reference systems as becomes evident if we look at the many spatial representation systems and calculi we can develop. All the different representations have advantages and disadvantages, depending on the problems we want to tackle or the situations we want to describe. Some problems can be solved directly on the object level; others are facilitated by suitable abstractions.

Nevertheless, spatial structures—and to a similar extent temporal structures—play special roles in everyday actions and problem solving. Many other dimensions seem to dominate our lives: monetary values, quality assessments, efficiency criteria, emotions, social structures, etc.—but do they play comparable roles with respect to cognitive representation and processing? I do not think so. I propose that the special role of space and time has to do with the fact, that internal representations may be a-modal, but they cannot be "a-structural". In other words: cognitive representations and processes depend on a spatio-temporal substrate; without such a substrate, they cannot exist. But they may not depend on a specific spatio-temporal substrate: a multitude of structures may do the job. Different abstractions from physical space may be advantageous in different situations.

Space and time provide fundamental structures for many tasks that cognitive agents must perform and for many aspects of the world that they can reason about. Maintaining these structures as a foundation simplifies cognitive tasks tremendously, including perceiving, memorizing, retrieving, reasoning, and acting. This is well known from everyday experiences, such as using geographic maps for wayfinding. For other domains it is helpful to create spatially structured foundations to support and simplify orientation; for example, spatial structure is the basis for diagrams that help us reason about many spatial and non-spatial domains.

A conceptually simple implementation of a truly spatial computer could be a robot system that manipulates physical objects in a spatial domain and perceives and represents these objects, the configurations constructed from these objects, and the parts of the objects as well as their relations from various orientations and perspectives. A more sophisticated approach would involve the construction of a spatial working memory, perhaps visually accessed, whose basic entities are entire objects rather than their constituents. Spatial operations like translation, rotation, and distortion would globally modify configurations. Perception operators extract qualitative spatial relations from these representations. The development of this implementation can be guided by our knowledge about working memory capabilities and limitations as well as by our knowledge about spatial representations in the human mind (Schultheis and Barkowsky 2011).

As technological materials become more sophisticated, the connection between spatial substrates and digital computer technology will become successively stronger. Sensor technology will be integrated into spatially structured materials in the years to come. A vision of such materials of the future that would support the concept of spatial computing on the substrate level can be found in a recent special issue in sensors and actuators (Lang et al. 2011).

12 A Final Note

Although we talk about spatial cognition, spatial reasoning, and spatial computing, we frequently fail to characterize the type of solution to spatial problems that we want to achieve. Our repertoire of approaches yields results on different levels of sophistication: some approaches only yield solutions to spatial problems; others yield some sort of explanations along with the solutions or instead of a solution. Accompanying explanations may be: ‘this is the only solution’; ‘this is one of possibly several solutions’; ‘these are all solutions’; or ‘there is no solution’.

Why is sophistication an issue? For highly abstract, formal approaches the quality of a solution is not obvious. Formal proofs or explanations (or both) are required to characterize the type of solution. In the more concrete, spatially structured solutions, the results are more easily perceptible, more obvious in that proofs may not be required—cognition and commonsense reasoning seem to operate without formal proofs, for the most part. On the other hand, can we be *sure* that we found the best solutions, the only solution, or all solutions? This is an old debate that calls into mind the discussion on the validity of constructive geometry to find solutions or to prove correctness.

There are different domains in which we can ground our knowledge: perceptual experience about spatial and temporal environments that does not require proofs and formal logics that does not require empirical justification. Both domains are important for human intellect and human reasoning. It does not make much sense to say one is superior over the other; they are two rather different realms. They may become particularly powerful when they are engaged jointly, one to carry out spatio-temporal perception and action and the other to reason about them on the meta-level and to explain what is going on in an overarching theory.

It is interesting to note that artificial intelligence research on commonsense reasoning so far has been restricted to characterizing commonsense reasoning on a descriptive level. Almost no AI work exists that emulates spatial cognitive abilities in a similar way as constructive geometry reflects spatial laws in the physical world or as artificial neural networks reflect topological structures akin to those of biological systems.

Acknowledgments I am grateful for valuable and detailed comments on earlier versions of this chapter from Thomas Barkowsky, Mehul Bhatt, Stefano Borgo, Holger Schultheis, Thora Tenbrink, Diedrich Wolter, several anonymous reviewers, and the editors of this book. Generous support from the German Research Foundation to the Spatial Cognition Research Center SFB/TR 8 Bremen and Freiburg is gratefully acknowledged.

References

Allen JF (1983) Maintaining knowledge about temporal intervals. CACM 26(11):832–843

Balbiani P, Condotta J-F, Ligozat G (2000) Reasoning about generalized intervals; preconvex relations and tractability. In: Proceedings of TIME-2000 conference, pp 23–30

Bhatt M, Lee JH, Schultz C (2011) CLP(QS): a declarative spatial reasoning framework. In: Egenhofer MJ, Giudice NA, Moratz R, Worboys MF (eds) COSIT'11, LNCS 6899. Springer, Heidelberg, pp 210–230

Cohn AG, Hazarika SM (2001) Qualitative spatial representation and reasoning: an overview. Fundamenta Informaticae 43:2–32

Condotta JF, Saade M, Ligozat G (2006) A generic toolkit for n-ary qualitative temporal and spatial calculi. In: TIME'06. IEEE Computer Society, Washington, pp 78–86 ISBN 0-7695-2617-9

DeHon A, Giavitto J-L, Gruau F (2007) 06361 Executive report — computing media languages for space-oriented computation. In: DeHon A, Giavitto J-L, Gruau F (eds) 06361 Abstracts collection—computing media languages for space-oriented computation, Schloss Dagstuhl. http://drops.dagstuhl.de/opus/volltexte/2007/1026

Edelman S, Intrator N (2000) Coarse coding of shape fragments + retinotopy ∼ representation of structure. Spat Vis 13:255–264

Egenhofer MJ, Franzosa RD (1991) Point set topological relations. Int J Geogr Inf Syst 5: 161–174

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

Freksa C (1991b) Qualitative spatial reasoning. In: Mark DM, Frank AU (eds) Cognitive and linguistic aspects of geographic space. Kluwer, Dordrecht, pp 361–372

Freksa C (1992a) Temporal reasoning based on semi-intervals. Artif Intell 54:199–227

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

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

Gentner D (1983) Structure-mapping: a theoretical framework for analogy. Cogn Sci 7:155–170

Greenwold S (2003) Spatial computing. Manuscript submitted to school of architecture and planning, Massachusetts Institute of Technology, Cambridge. http://acg.media.mit.edu/people/simong/. Accessed 1 Jan 2012

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

Hamblin CL (1972) Instants and intervals. In: Fraser JT, Haber FC, Müller GH (eds) The study of time. Springer, Berlin, pp 324–331

Hwang A (2012) http://albert-hwang.com/projects/spatial-computing/. Accessed 1 Jan 2012)

Kosslyn SM (1987) Seeing and imagining in the cerebral hemispheres: a computational approach. Psychol Rev 94:148–175

Lang W, Lehmhus D, van der Zwaag S, Dorey R (2011) Sensorial materials—a vision about where progress in sensor integration may lead to. Sens Actuators, A 171(1):1–2

Ligozat G (1993) Qualitative triangulation for spatial reasoning. In: Campari I, Frank AU (eds) COSIT 1993 LNCS 716. Springer, Heidelberg, pp 54–68

Liu W, Li S (2011) Reasoning about cardinal directions between extended objects: the NP-hardness result. Artif Intell 175:2155–2169

Moratz R, Renz J, Wolter D (2000) Qualitative spatial reasoning about line segments. In: Proceedings of ECAI 2000, pp 234–238

Moratz R, Tenbrink T, Fischer F, Bateman J (2003) Spatial knowledge representation for human-robot interaction. In: Freksa C, Brauer W, Habel C, Wender KF (eds) Spatial cognition III – Routes and navigation, human memory and learning, spatial representation and spatial reasoning, LNAI 2685. Springer, Heidelberg, pp 263–286

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

Newell A (1982) The knowledge level. Artif Intell 18(1):87–127

Nicod J (1924) Geometry in the sensible world. Doctoral thesis, Sorbonne, English translation in Geometry and Induction, Routledge and Kegan Paul, 1969

Randell DA, Cui Z, Cohn AG (1992) A spatial logic based on regions and connection. In: Proceedings of 3rd international conference on knowledge representation and reasoning. Morgan Kaufman, Los Altos, pp 55–66

Rosch E (1978) Principles of categorization. In: Rosch E, Lloyd BB (eds) Cognition and categorization. Erlbaum, Hillsdale

Schultheis H, Barkowsky T (2011) Casimir: an architecture for mental spatial knowledge processing. Top Cogn Sci 3:778–795

Schwartz JT, Sharir M (1983) On the "piano movers" problem I: the case of a two-dimensional rigid polygonal body moving amidst polygonal barriers. Commun Pure Appl Math 36: 345–398

Shepard RN, Metzler J (1971) Mental rotation of three-dimensional objects. Science 171: 701–703

Skiadopoulos S, Koubarakis M (2004) Composing cardinal direction relations. Artif Intell 152:143–171

Sloman A (1971) Interactions between philosophy and artificial intelligence: the role of intuition and non-logical reasoning in intelligence. Artif Intell 2:209–225

Sloman A (1985) Why we need many knowledge representation formalisms. In: Bramer M (ed) Research and development in expert systems. Proceedings of BCS expert systems conference 1984. Cambridge University Press, Cambridge, pp 163–183

Tobler W (1970) A computer movie simulating urban growth in the Detroit region. Econ Geogr 46(2):234–240

Van de Weghe N, Kuijpers B, Bogaert P, De Maeyer P (2005) A qualitative trajectory calculus and the composition of its relations. In: Proceedings of geospatial semantics, vol 3799, pp 60–76

Wallgruen JO, Frommberger L, Wolter D, Dylla F, Freksa C (2007) Qualitative spatial representation and reasoning in the SparQ toolbox. In: Barkowsky T, Knauff M, Ligozat G, Montello D (eds.) Spatial cognition V: reasoning, action, interaction, LNAI 4387. Springer, Heidelberg, pp 39–58

Westphal M, Woelfl S, Gantner Z (2009) GQR: a fast solver for binary qualitative constraint networks. In: AAAI spring symposium on benchmarking of qualitative spatial and temporal reasoning systems, Stanford

Wolter D, Wallgruen JO (2012) Qualitative spatial reasoning for applications: New challenges and the SparQ toolbox, In: Hazarika SM (ed) Qualitative spatio-temporal representation and reasoning: Trends and future directions. IGI Global, Hershey. doi: 10.4018/978-1-61692-868-1

Yang C, Goodchild M, Huang Q, Nebert D, Raskin R, Xu Y, Bambacus M, Fay D (2011) Spatial cloud computing: how can the geospatial sciences use and help shape cloud computing? Int J Digit Earth 4:4

Zimmermann K (1995). Measuring without measures. The Δ-Calculus. In: Frank AU, Kuhn W (eds) COSIT 1995, LNCS 988. Springer, Heidelberg, pp 59–67

Zuse K (1969) Rechnender Raum: Schriften zur Datenverarbeitung. Vieweg, Braunschweig