

Spatial Problem Solving and Cognition

Christian Freksa, Thomas Barkowsky, Frank Dylla, Zoe Falomir,
Ana-Maria Oltețeanu, Jasper van de Ven

Bremen Spatial Cognition Center, University of Bremen, Germany

1 Spatial problems and physical space

A spatial problem is (1) a question about a given spatial configuration (of arbitrary physical entities) that needs to be answered (e.g. *is there wine in the glass?*) or (2) the challenge to construct a spatial configuration with certain properties from a given spatial configuration (e.g. *add two matchsticks to the given configuration to obtain four squares*) (Bertel, 2010). By spatial configurations we mean arrangements of entities in 1-, 2-, or 3-dimensional physical space, where physical space is commonsensically observable Euclidean space and motion, rather than relativistic space-time. Physical space is contrasted here to abstract space of arbitrary dimensionality. Physical space affords certain actions, like (i) rotation (circular motion of objects around a given location); (ii) motion from one location to another; (iii) deformation of objects; (iv) separation of objects into parts; (v) aggregation of objects; and (vi) combinations, i.e. rotation around a changing location.

A special feature of *commonsense physical space* (CPS) is that operations such as motion are severely constrained and comply with rigid rules we cannot change whereas in abstract spaces we are free to make up arbitrary rules about which operations are possible and which are not. For example, in *abstract representations of space* (AbsRS) we could allow a ‘jump’ operation that moves an entity directly from one location to a remote location (as in some board games). In CPS this is not possible: objects always first move to neighboring locations and then to a neighbor of that location, etc., before they can reach a remote location¹. This has implications on the trajectories (including the time course) of motion.

As a second example, in abstract space we could come up with an operation that allows an entity to be in two places at the same time. In CPS this is not possible because of the nature of physical space and matter. This has implications on unique identity, presence in a space, containment within it and access to it. In abstract space, the types of operations possible are defined by the agent conceiving the abstract space, while in CPS they depend on the nature of physical space itself. The types of actions that can be performed in CPS define the characteristic structure of physical space (Freksa, 1997) that is exploited by Euclidean geometry and vice versa (Euclid, 300 BC/1956).

In this chapter, we discuss (1) how cognitive agents such as humans, other animals, or robots can use concrete CPS and AbsRS for solving spatial problems and (2) what are the relative merits of both approaches. We describe how the approaches can be combined. We look at the roles of spatial configurations and of cognitive agents in the process of spatial problem solving from a cognitive architecture perspective. In particular, we discuss (a) the role of the structures of space and time; (b) the role of conceptualizations

¹ This holds for arbitrary granularities of neighborhoods.

and representations of these structures; and (c) the role of knowledge about these structures.

The chapter is organized as follows. In section 1 we describe how and why geographic maps help us solve spatial problems in the real world (1.1); then point out cognitive difficulties of communicating about space and spatial representations (1.2); and offer a wayfinding example that illustrates how various levels of abstraction can be involved in spatial problem solving and reasoning. A fresh look at spatial problem solving is taken in section 2, where we describe components of problem solving (2.1); put the components together (2.2); and discuss the difference between *solving* spatial problems and *understanding* problem solving processes (2.3). On the basis of this discussion, section 3 proposes mild abstraction as a third way between direct spatial and indirect formal problem solving, discusses how much abstraction is useful (3.1); illustrates how mild abstraction is performed in geographic maps (3.2); moves the discussion of mild abstraction from geographic space to other spatial domains (3.3); and discusses strategic aspects of applying this approach to spatial problem solving (3.4). Finally, section 4 concludes with a discussion of three levels of cognitive processing in spatial problem solving and implications for cognitive approaches to spatial problem solving.

1.1 Physical representation of space: Geographic maps

Maps distort the space in which we want to navigate. Most notably, maps shrink the space to such an extent that it is not possible to walk or drive in map space. Or stated differently: although maps represent space physically by means of a concrete spatial medium², they abstract from certain aspects of the spaces they represent. Specifically, they systematically substitute distances by smaller distances in such a way, that certain other aspects of space (e.g. connectivity, orientation, angles, relative distances, or relative areas) are preserved.

Thus, spatial relations in environmental space are projected into similar spatial relations in map space. Due to the specific analogical way of representing spatial relations by identical or projections of spatial relations (Sloman, 1971; Robinson et al., 1995) we perceive these spatial relations in the map as if we would perceive them in the environment under more favorable perception conditions (more suitable perspective, scale adapted to our field of view, no obstructions) (MacEachren, 1995). As far as spatial relations are concerned, it is as if we would look at the spatial environment with a demagnifying lens or from a large distance above the ground. If we have the map at our current location, we can do this without the effort of moving about the environment.

The map offers a bird's eye view of an ample set of spatial relations that we could rarely observe all at once in the environment, without the help of a high point, like the peak of a mountain or a hot air balloon. Allowing us to perceive many more spatial relations at once, rather than keep in mind some objects while we move around and discover others, and then establish the relation, the map thus acts as an extended memory with visual access, and supports a form of external cognition (Scaife & Rogers, 1996; Tversky & Lee, 1999; Card et al., 1999). The operations we can perform on the map are basically the same as the ones we can perform in the environment (e.g. triangulation, measuring distances, path following); however, we must be aware of distortions if we project from

² We will use 'representation' and 'medium' synonymously in this chapter. While 'representation' emphasizes structural aspects, 'medium' emphasizes physical and spatial aspects.

a sphere to the flat surface of a map (Monmonier, 1996). As geographic maps *preserve essential spatial relations* implicitly *as spatial relations* and do not abstract them away, we consider map representations as *mild abstractions* of the spatial environment they represent.

Barbara Tversky early on pointed out that *cognitive* maps are spatially distorted wrt. the represented space (Tversky, 1981, 1992, 1993; Mark et al., 1999); this insight inspired the spatial cognition community including the present authors to investigate potential advantages of spatial distortions. For example, why is it easier to navigate with distorted subway maps than with veridical maps even though it should be more difficult to match those maps to the environment (Berendt et al., 1998)? Once we recognize that (physical or cognitive) spatial distortions actually may simplify spatial problem solving, we open up a whole new domain for studying spatial problem solving.

1.2 On the difficulty of communicating about space and spatial representations

A cognitive issue that causes problems when discussing spatial representations and spatial cognition is the following: in human language we often do not distinguish between entities in the real world and their physical or mental representation. For example, we point with a finger to a location on a map and explain to another person *now we are here*; we intend to express that we are located at the place in the environment that is represented by the corresponding location on the map that we are pointing at.

An interesting aspect that contributes to the confusion between the environment and its representation is that, really, we do not care so much about where we are *on the map*; we are in fact interested to know where we are in the environment. The answer to the latter however would be trivial and not useful: we are *right here* where we are standing; this answer even is highly context-adaptive, i.e. it is valid wherever we are!

Interestingly, it is frequently easier to find out in a completely different space (map space) where we are, than in the environmental space itself. The reason is that, on the map, our perception provides us with an overview of a multitude of known locations (pardon: representatives of known locations) (Tversky, 2000); in this way, we are able to relate the representation of our location to the representations of the locations of other entities, and thus perceive relations (in front, right, left, north, south, etc.) between these representations. These representations allow us to derive relations between the corresponding locations in the environment. It is through these relations that we are able to orient ourselves and understand our environment, in the same way in which a listener requires a few bars of music to pass before they can tell in which tonality they are “located” - finding their “place” through a web of musical relations.

Therefore, for orientation purposes, there is no necessity to distinguish between the environment and its map representation: the map serves as an aid to perceive the environment and may equally well be considered a part of our perception apparatus (a de-magnifying lens) as it can be viewed as a space that is conceptually outside the geographic environment.

When we discuss cognitive processes as scientists, we have a different situation than when we try to locate ourselves: as scientists, we need to carefully distinguish the spatial environment from its representation. But in practice, the problem domain and the representation domain are conflated even in scientific contexts as if they were

identical. For example, in much of the artificial intelligence (AI) work, spatial problems are defined on the formal representation level, where – unlike in our map example – no relevant spatial relations are intrinsically given (Russell & Norvig, 1995). On the formal level, however, we are dealing with *descriptions of* spatial relations rather than with spatial relations. Formally trained people read these descriptions as if they were the spatial relations themselves, just as trained map-readers read maps as if they were the environments themselves.

Descriptions of spatial relations make some properties explicit which are implicitly present in physical spatial structures; they convert these properties into *knowledge about* the properties. For example, a distance between two cities is implicitly given through the distance between the cities' locations on the map; knowledge about this distance could be expressed explicitly for example by “distance (city_A, city_B) = 40 km” or by specifying the cities' coordinates and computing the distance through an explicitly specified algorithm.

Symbol systems in AI (and formally trained people) use these descriptions to *reason about* spatial relations. This allows solving spatial problems indirectly, by arguing about what effects spatial relations and properties of spatial structures would have if we were to solve a real spatial problem. Some AI researchers seem to suggest that cognitive agents including humans and other animals must solve spatial problems by reasoning about them (Davis, 2013); but do toddlers or dogs reason about spatial relations when they open a door? How can cognitive agents solve spatial problems if they are lacking the explicit symbolic knowledge needed for reasoning? Explanations are rarely given.

1.3 A wayfinding example: Finding a shortest path between two locations

To appreciate some of the issues involved in spatial problem solving and reasoning, let us consider approaches to determine a shortest path between two locations in a route network. We will first sketch how this problem could be solved directly in the spatial environment; we will then discuss different ways of solving this problem with the support of various kinds of representations.

1.3.1 Finding a shortest path in the spatial environment

To determine a shortest path in a route network, we must (a) be able to compare lengths of paths and determine which of two paths is shorter; and (b) take into account all possible paths between the start and the end points of the respective route network, to be sure we identified a shortest path.

Unless we can directly relate and perceive the extent of two paths, it is difficult to compare their lengths, as we lack sensors to compare path lengths; therefore we have to resort to some indirect way of comparing lengths: for example, we can identify the start point of a route with one end of a rope that we stretch out along the route; we can mark the rope at the end of the path (provided it is long enough); we then can move the marked rope section to another path and determine whether its length is less, equal, or more than the marked segment of the rope. In doing so, we assume that the length of ropes is preserved when they are moved and we make use of the transitivity property of length: if the length of the rope section is the same as the first path, then comparing the second path to the rope yields the same result as comparing the second path to the first

path would yield if we could do it. Note that we do not have to measure lengths quantitatively in order to compare them.

There are other methods for indirectly comparing lengths in spatial environments; popular ones include: counting the number of steps of constant length and comparing the step counts; moving along the paths at constant speed and comparing the traversal times (this can be done qualitatively by comparing contents of sand clocks or quantitatively by measuring times that then can be compared).

An additional challenge will be to make sure that we take into account all possible paths; this requires an ability to identify paths and to record whether they have been compared to another path, and if so, to which.

In summary, if we solve the shortest path problem directly in the spatial environment, we will require tools for comparing lengths and for keeping track of the problem-solving progress. These tools can be part of the spatial environment. A perceiving agent is required to assess differences (in lengths and path identity).

Solving the shortest path problem directly in the spatial environment is cumbersome. Largely this is due to the size of the environment and our lack of sensors that can cope with this size; therefore we scale down the size. As we are interested in the role the representation medium or the structure of the representation plays for the problem solving process, we will first consider various media that we can use to solve navigation problems: a map or visual graph; an abstract graph; and a list.

1.3.2 Finding a shortest path in a map

Let us suppose you are using a map like that in Fig. 1 and you want to find a shortest path from the intersection of Normandie Ave. and W 35th Pl (symbol ⊗ on the map) to University of Southern California (symbol ⊙ on the map). In order to find a shortest path, one could naively apply the same methods as in the ‘spatial environment’, i.e. comparing all possible routes. One advantage of a map representation is the provision of overview knowledge. The straight connection between start position and destination provides a direct means to compare the currently considered path to a shortest possible connection, the linear distance. If there are several candidates whose length cannot be discriminated visually, one may have to compare them by some indirect means, e.g. by using a string in the same way as the rope in the ‘spatial environment’ example (section 1.3.1).

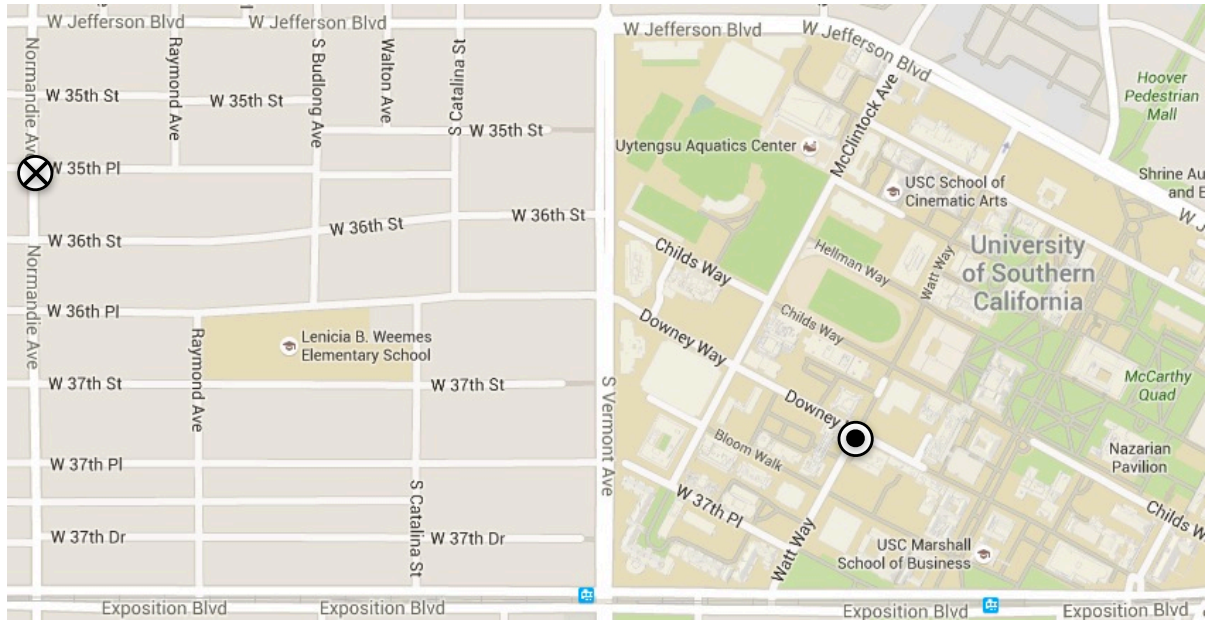


Figure 1. Section of a city map for solving a wayfinding problem with visual and haptic support.

Task-irrelevant information, like buildings or parks, does not have to be represented. In the abstraction of the spatial environment that leads to the map, no task-relevant information is lost, except perhaps altitude information or some other spatial distortion, depending on the specific map projection employed. The resulting schematic map (Fig. 2) is spatially equivalent to the map as long as spatial relations are preserved, i.e. the graph provides the same task-relevant spatial information.

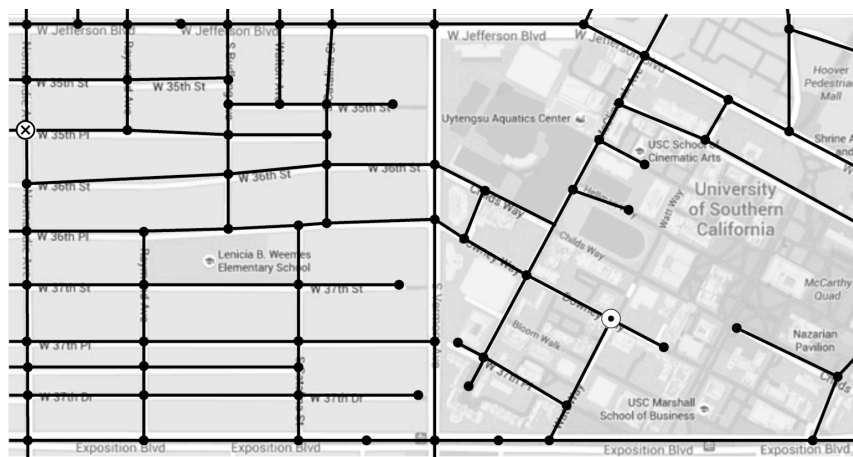


Figure 2: Distance-preserving schematic map providing spatial information for path finding – superimposed on the map.

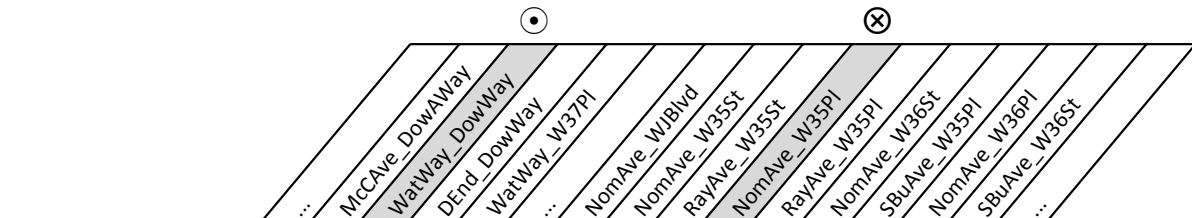
We can go one step further to facilitate the task of finding a shortest path: we can construct the schematic map as a ‘string map’ from flexible non-elastic strings, that connect the nodes of the spatial graph and preserve the lengths of the edges. A shortest path then can be found by a simple physical action: we pull apart the nodes that correspond to the starting position and the target position until we obtain a straight

connection between them; the strings on the taut connection represent a shortest path (Dreyfus & Haugeland, 1974; Freksa et al., 2016).

1.3.3 Finding a shortest path in an abstract graph structure

The next representation we consider is an abstract graph. Spatial information is provided by means of vertices representing road junctions and edges representing road connections. Abstract graphs do not convey distance information of paths implicitly as maps do; therefore, we label edges explicitly with numerical values that reflect the distance between junctions along the corresponding path. With the abstraction from a map, the connection between the spatial environment and the representation is lost. Although this connection is not relevant for solving the abstract shortest path problem, it will be required to apply the abstract solution to the real world. This can be achieved by explicitly annotating abstract graphs with coordinates or names of locations.

Our abstract graphs focus on road junctions and connections between them. Apart from that, space is not represented. Abstract graphs can be coded in various ways. A popular representation scheme for graphs is an adjacency matrix. In our case, the matrix will contain one column and one row for each vertex in the graph (Fig. 3). The elements of the matrix indicate whether the pairs of vertices are adjacent or not in the graph. As we want to identify shortest paths, we will enter the length of the corresponding path segment in the matrix and leave the entries for non-adjacent vertex pairs empty. This will enable suitable computer algorithms to determine accumulated path lengths of chains of path segments. Note that – unlike in the map where each location in the environment is represented by a unique location – each vertex is represented twice in the adjacency matrix: once as a starting node (row) and once as an ending node (column) of a connection.



	...	McCave_DowAWay	WatWay_DowWay	DEnd_DowWay	WatWay_W37Pl	...	NomAve_WJBlvd	NomAve_W35St	RayAve_W35St	NomAve_W35Pl	RayAve_W35Pl	NomAve_W36St	SBUAve_W35Pl	NomAve_W36Pl	SBUAve_W36St	...
...
● McCave_DowAWay	...	0	32													
● WatWay_DowWay	...	32	0	12	27											
DEnd_DowWay	...		12	0												
WatWay_W37Pl	...		27		0											
...
NomAve_WJBlvd	...					0	15									
NomAve_W35St	...					15	0	33	16							
RayAve_W35St	...					33	0		15							
⊗ NomAve_W35Pl	...					16		0	32	15						
RayAve_W35Pl	...						15	32	0		32					
NomAve_W36St	...							15		0		16	65			
SBUAve_W35Pl	...								32		0					
NomAve_W36Pl	...									16		0				
SBUAve_W36St	...									65			0			
...

Figure 3: Adjacency matrix representing a segment of the visual graph shown in Fig. 2.

Fig. 3 also uses a (2D) spatial medium for the representation; but the space of the medium no longer carries spatial information about the environment. The *spatiality of the medium*, however, still facilitates our perception of the connection relations between the vertices, as it provides an easy overview of these relations to our visual system at a glance. Computer algorithms usually do not make use of the *spatiality of representations*; they ‘look up’ one connection after another and construct abstract chains of the route segments and accumulate their length values to determine a shortest path.

Fig. 4 visualizes the abstract graph represented in Fig. 3 in a more human-friendly fashion. As the spatial medium no longer carries spatial information, this visualization is informationally equivalent to Fig. 3.

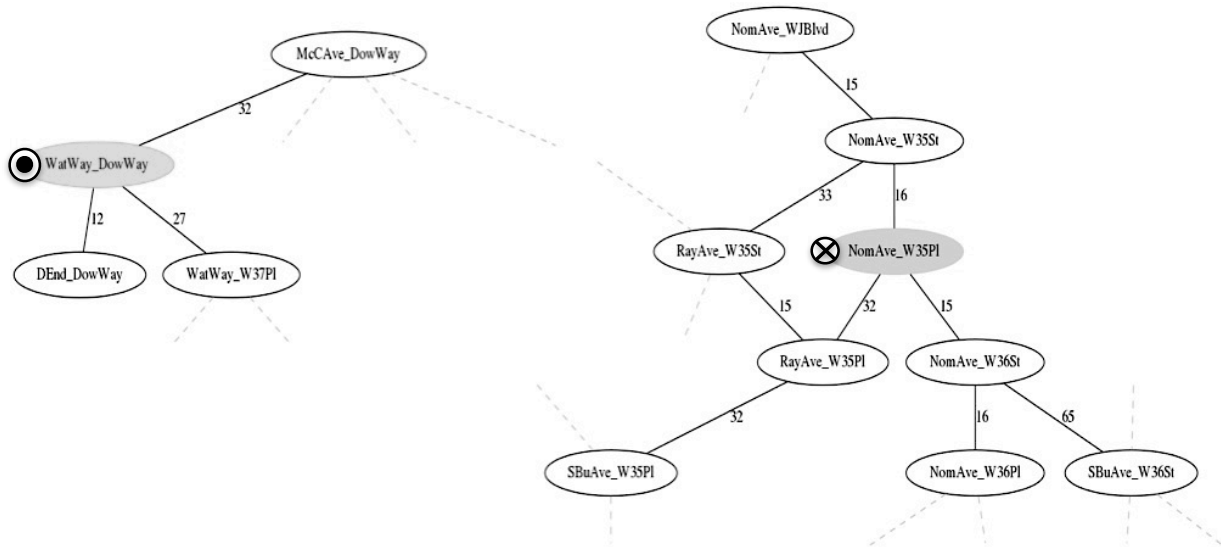


Figure 4: Visualization of the abstract graph represented in Fig. 3. Vertices depict road junctions and edges depict road connections; the labels indicate lengths of route segments.

The graph visualization in Fig. 5 depicts the same road network as Fig. 1. This depiction is just one possible visualization, which does not necessarily reflect spatial relations in the spatial environment. An arbitrary number of different visualizations can be generated from an abstract graph. This implies that the spatial methods to determine a shortest path, as introduced in sections 1.3.1 and 1.3.2, are not applicable to this kind of representation. For comparing lengths, we can no longer compare routes visually; instead, we must interpret and compare numerical values. The overall path length is determined by the accumulation of the lengths of the individual segments. As a consequence, we must compare all possible paths, e.g. by means of an algorithm like the Dijkstra algorithm (Dijkstra, 1959).

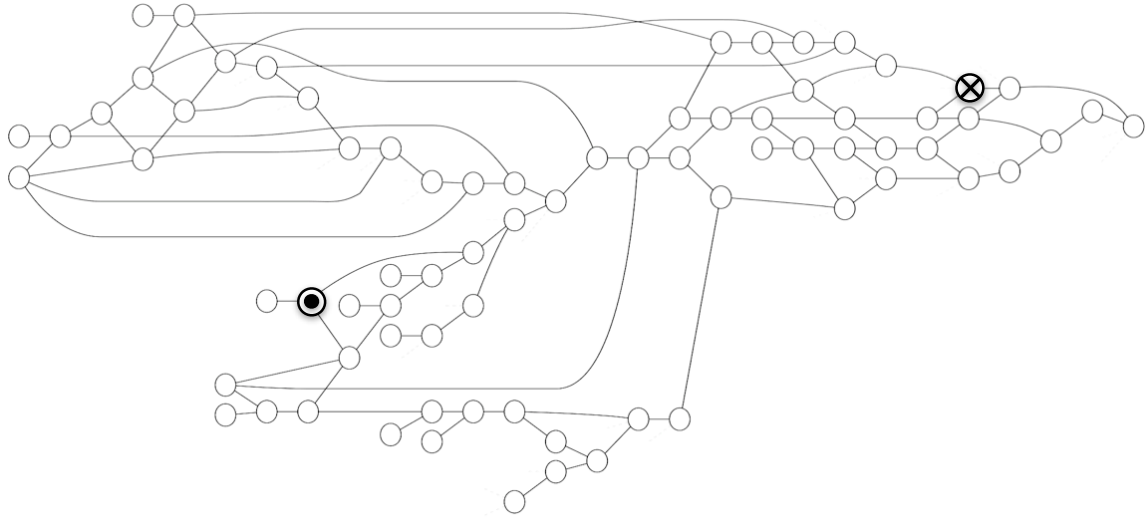


Figure 5: One possible visualization of the complete abstract graph representation of the road map shown in Fig. 1 without vertex names and distance labels.

1.3.4 Finding a shortest path in a list of path segments

Adjacency matrices (Fig. 3) for route connections typically comprise a large fraction of empty entries, as vertices typically are directly connected only to some of the other vertices. As a consequence, we can compress the relevant information without losing task-relevant information by focusing on the edges between the vertices and explicitly representing only those relations between vertices that comprise a direct connection. To this end, the information about the edges can be represented by a list of triples that each contain two labels representing the vertices and an associated path length: (<start-of-path-segment>, <end-of-path-segment>, <length-of-path-segment>).

For example, we may have a list that contains the following triples (cf. Fig. 3):

```
(WatWay_DowWay, McCave_DowWay, 32)
(DEnd_DowWay, WatWay_DowWay, 12)
(NomAve_W35PI, RayAve_W35PI, 32)
(RayAve_W35PI, SBuAve_W35PI, 32)
...
```

The list requires only as many entries as there are direct paths between vertices. Each vertex appears as many times in the list as it functions as a starting or end point of a connection. The order of the triples in the list is insignificant (i.e. the list is interpreted as a set; no information is conveyed by the sequence of its elements).

In this representation, the graph has been chopped into pieces; spatial integrity is lost. However, all the information required to reconstruct the graph correctly has been maintained; thus, the graph can be reconstructed computationally by linking triples that comprise identical vertex names. Consequently, shortest path finding algorithms like the Dijkstra algorithm can be applied.

1.3.5 Summary

Table 1 summarizes the progressive transformation from spatially implicit information in environments and maps to symbolically explicit information about spatial paths in

abstract graphs and lists of edges. The mild abstraction of spatial information in the environment into veridical or schematized maps maintains task-relevant essential spatial features and adds perceptual, haptic, and mental affordances for human use, by adapting to the visual and haptic field of humans and permitting the use of global path-finding heuristics.

The transformation from maps to abstract graphs switches from a space-based representation with implicitly maintained spatial relations to a feature-based representation with symbolically explicit representation of spatial features. In our examples, the abstract graph is represented by a junction-based adjacency matrix that makes limited use of structural landmark association by means of rows and columns to represent information implicitly. Most other task-relevant spatial features are made symbolically explicit.

The transformation from abstract graphs to a list of edges that represent direct path connections serves a compaction of the task-relevant information and no longer makes use of spatial / positional information (except within the triples that denote the edges). In the transformed representations, some information is implicitly assumed that permits to reconstruct spatial integrity to some extent. For example, we assume that identical location identifiers refer to a unique location (and in some systems we assume that distinct location identifiers refer to distinct locations); this permits us to reconstruct correctly connected graph structures from rather sparse information about edges.

Table 1: Progressive transformation from spatially implicit information in spatial environments and maps (left) to symbolically explicit information about paths in abstract graphs and lists (right).

	Environment	Map	Abstract graph	List of edges
Spatial abstraction	(none)	Mild abstraction: Space-based => space-based <ul style="list-style-type: none"> Relative locations preserved Scale: absolute distances => relative distances Dimensions: 3D => 2D 2D connectivity preserved 2D orientations preserved Spatial integrity is largely maintained	Transformation: Space-based => junction-based <ul style="list-style-type: none"> Locations of junctions: implicitly unique => conceptually separated into start & end of multiple path segments Distances: implicit => explicit Path segment connectivity: implicit through matrix structure and labels Spatial integrity is partially dissolved	Compaction: Junction-based => path segment-based <ul style="list-style-type: none"> Location of junctions: start & end of multiple path segments => start & end of individual path segments Distances: explicit Path segment connectivity: implicit through label uniqueness Spatial integrity is fully dissolved
Spatial features	(all present)	<ul style="list-style-type: none"> Absolute 2D location explicit (map section) Relative 2D locations implicit (spatial medium) Spatial scale (explicit / implicit) Connectivity relations (spatially implicit) 2D orientations (spatially implicit) Connectivity between path segments implicit Implicit representation of missing connections 	<ul style="list-style-type: none"> Locations of junctions (explicit by label) Direct distances between junctions (explicit by label) Junction identity (partly spatially implicit in matrix; partly symbolically implicit through label) Direct connections between junctions explicit Connectivity between path segments implicit Explicit representation of missing direct connections 	<ul style="list-style-type: none"> Locations of junctions (explicit by label) Direct distances between junctions (explicit by label) Junction identity (symbolically implicit through label) Direct connections between junctions explicit Connectivity between path segments implicit Implicit representation of missing direct connections
Spatial affordances	<ul style="list-style-type: none"> Physical motion through environment until goal is reached Random path selection 	<ul style="list-style-type: none"> Overview perspective Perceptual path length comparison Orientation-based heuristic path selection Haptic or mental simulation of path traversal 	<ul style="list-style-type: none"> Partial junction identity due to matrix structure 	(none)

The spatial integrity that collocates all features of a spatial location at that location in a spatial environment is progressively dissolved in the transitions to map, graph, and list respectively, such that all spatially implicit information in the environment that is represented will be made symbolically explicit at the final stage.

2 A fresh look at spatial problem solving

Many spatial problems are solved every day *without representing* them as spatial problems in the mind or the computer. For example, my keys open locks mechanically without a representation of the lock's mechanism needing to exist in my mind; doors open by leaning against them or “magically” through sensor-controlled mechanisms that respond to my approaching the door. In these cases, physical affordances (Gibson, 1979) established in the interaction between the environment and the agent enable solutions to spatial problems without reasoning needing to be involved.

While this type of problem solving may be intellectually unsatisfactory for computer scientists, as cognitive scientists we must acknowledge that such an action and perception-based approach developmentally precedes reflective thinking and most likely is a prerequisite for building up mental representations of spatial problems and for spatial problem solving (Johnson, 2009; Needham, 2009; Keen, 2003). Initially, cognitive agents can relate cause and effect of actions; later, they can describe and possibly understand the underlying process. At that point, cognitive agents have a representation that may enable them to find problem solutions mentally or computationally by reasoning. Once they have generated a solution by reasoning, they can apply it to the actual spatial situation in the environment.

We can distinguish two types of processes involved in spatial problem solving: problem solving processes that operate in a given medium such as the physical space, a geographic map, a logic formalism, or some other representation of space; and problem transformation processes that transform problems between different media or kinds of representations. AI problem solving has been largely concerned with the first type of process (e.g. Fikes & Nilsson, 1971); but there also are approaches on the formal level that re-represent a given problem in a different formalism in order to determine a problem solution more easily (Yan et al., 2003). The role of different representations has been discussed by Bobrow (1975), Palmer (1978), Marr (1982), Sloman (1985), and Freksa (2015); the importance of paying attention to the transformation between media or forms of organizing knowledge to achieve such transformations were demonstrated by Freksa (1988) and Oltețeanu (2016).

Different media (physical space, map space, diagrams, logic representation, etc.) afford different operations and thus favor different kinds of problem solving. Therefore, solving the same problem in different media or representations results in different process structures and possibly in a different scalability of the problem solving process (Larkin & Simon, 1987).

2.1 Components of spatial problem solving

Spatial problems can be given in physical space or in an abstract form; similarly, the problem solution can be given as a spatial configuration or in the form of an abstract description. Accordingly, we can distinguish different kinds of spatial problem solving processes, depending on whether they are performed entirely within the spatial domain, entirely within the abstract domain, or by some sort of a combination.

For the present discussion, we will focus on problems that are given in physical space and for which the solution sought is a spatial configuration. For this situation, we can identify four basic spatial problem-solving components that may be involved in these processes:

- 1 Solving a physically spatial problem by operating on the problem configuration *within* the spatial medium
- 2 Solving a spatial problem by processing the description of the problem
- 3 Transforming a physically-spatial problem into a differently structured (e.g. abstract) representation medium
- 4 Transforming the description of a problem solution from the representation medium to a spatial configuration

These components and their interrelationships are depicted in Fig. 6.

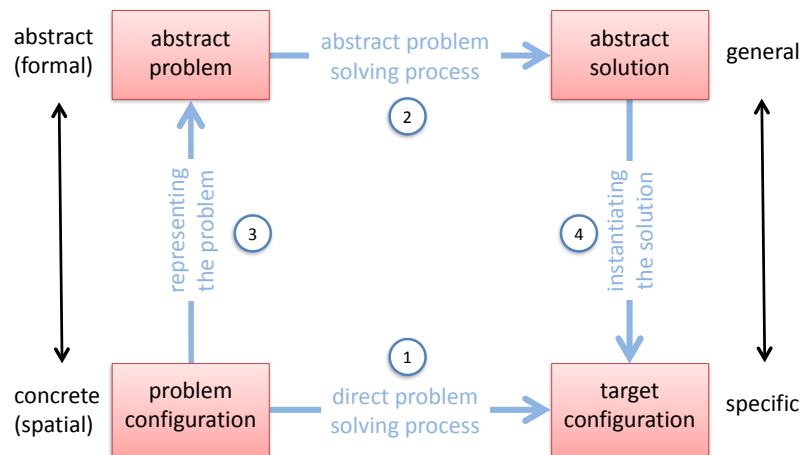


Figure 6. From spatial problem to spatial solution: A spatial problem configuration (bottom left) can be operated on directly in space to obtain a spatial target configuration (bottom right); alternatively, it can be transformed by abstraction into a mental or formal representation (top left), mentally or computationally processed into a solution (top right), which then can be re-transformed into a spatial configuration.

2.2 Putting the components together

Cognitive agents may get away with limiting their approach to process component 1 applied directly to the spatial medium: operating on the spatial problem configuration through a physical action, in order to obtain a spatial solution configuration that manifests the desired effect (e.g. opening a door by leaning on it). This can be achieved by trial and error that accidentally solves the problem (due to spatial affordances) or by an intentional action that has known effects, but has no explanation of how the effects are produced in terms of process.

Cognitive agents can solve spatial problems in a variety of ways: they can (A) take an action that solves the problem (i) accidentally or (ii) by means of known effects; in either case spatial affordances enable the problem solution process. Or they can (B) transform the spatial problem into a formal problem for which (i) a solution can be searched for on the formal level or for which (ii) a solution is already known. Or they can (C) transform the spatial problem into a non-formal representation such as a map or a diagram that (i) may facilitate visual search for a solution or for which (ii) a solution is already known (see section 3).

The fact that cognitive agents have such a variety of options for spatial problem solving at their disposal suggests that being capable of pursuing several of these options requires some sort of meta-knowledge (or intuition) regarding which direction to

pursue in order to solve a given problem. This meta-knowledge involves knowledge about the objectives to be achieved and knowledge or beliefs about strategies that may be successful. In the following section (2.3) we will address the issue of objectives; in section 3.3 we will address the issue of strategies.

2.3 Solving problems vs. understanding problem solving

Spatial problem solving may serve different objectives: we may solve a problem in order to obtain a desired spatial configuration (e.g. an open door to pass through) or we may investigate problem solving in order to *understand* spatial transformation processes and principles intellectually as scientists. Although these two objectives are related, they require different kinds of models: spatial reconfiguration requires physical action, and actions need to be controlled by some sort of a mind; **a mind that controls actions requires different knowledge than a mind that understands spatial transformations and their implications.** Understanding spatial transformations does not necessarily require physical action; it requires notions of causality, topological transformation, geometric equivalence, and logical inference.

Much of the research on spatial problem solving since Euclid has focused on the intellectual challenge of *understanding spatial structures and principles* underlying spatial operations, as well as their implications regarding spatial problem solving: it concerned general spatial problems that we wanted to solve as scientists and whose solutions indirectly also could solve a common agent's problem of how to transform specific spatial configurations into desired target configurations. Consequently, this research has been concerned with formal descriptions of space and its properties: a representation that has proven particularly useful for intellectual treatment and analysis.

Understanding the principles of physical space certainly is most useful for characterizing the abstract space of potential approaches to (concrete) spatial problem solving. But is an understanding of formal spatial principles and structures sufficient to replicate or synthesize the kind of *commonsense spatial problem solving* exercised by animals and common people (including toddlers) who lack the explicit knowledge to reason about spatial principles and structures?

We believe that the difference between solving problems and understanding their solution is underappreciated. This is partly due to the mental identification of spatial situations and their representation that was described in section 1.2. Although *embodiment* and *situatedness* have become elaborately described and generally accepted notions, we find little work that employs this approach to spatial problem solving. We are convinced that it is worthwhile investigating naïve spatial problem solving by studying spatial affordances generated in the interaction between cognitive agents and spatial configurations. Our aim is to understand and synthesize naïve cognitive agents who depend on pre-intellectual shortcuts in spatial problem solving; these agents make direct use of spatial structures without knowing any principles that could guide their actions. Thus, our approach can be viewed as an action and perception-based approach that complements existing symbol-based approaches.

3 Mild abstraction: A third way between direct and formal problem solving

As we have seen, solving spatial problems directly in a spatial medium and solving problems abstractly both may have remarkable advantages, depending on the specific problem-solving requirements. We have seen that cognitive agents applying their perception and action capabilities to maps and visual graphs can make use of implicit spatial structures to solve spatial problems – why should we not equip computational problem-solving systems with the same type of capabilities? This will enable them to make direct use of multiple levels of features that are integrated into spatial structures. This also will provide shortcuts to purely computational spatial problem solving as it avoids spatial disintegration and re-integration steps involved in symbolic information processing.

The key idea is to implicitly maintain those aspects of the problem domain that can directly support the spatial problem solving process, and to formalize only those aspects that do not support a solution in the spatial domain, but can be used for explicit reasoning in the formal domain (Freksa, 1991; Freksa et al., 2000). A main advantage will be that, by manipulating spatial structures, we will simultaneously manipulate coarse and fine levels of space, as well as all aspects that are integrated in the spatial representation. Spatial integration thus also offers a solution to the frame problem (McCarthy & Hayes, 1969).

3.1 How much abstraction is useful?

A great advantage of using spatial structures for spatial problem solving is that crucial properties of space, such as topological and geometric laws including dimensionality and inherent relations between dimensions, are maintained intrinsically and therefore can be directly exploited without any need to reason about them (Palmer, 1978; Dirlich et al., 1983; Freksa, 2015; Furbach et al., 2016). Generalization to a wider range of sizes, orientations, and geometries may be desirable; but generalization beyond the general constraints of space would not be useful when we want to use our representation exclusively to solve truly spatial problems. On the contrary: if we relax constraints generally applicable to spatial domains, we have to invest additional computational effort on the representation level in order to guarantee results that conform to the realm of space (Freksa, 1997).

On the other hand, generalization to less constrained representations may be quite useful for other kinds of problems, e.g. if we want to reason about abstract mathematical spaces or conceptual spaces (Gärdenfors, 2000), where we explicitly intend to escape the confinements of physical space.

But also for purely spatial problems there are aspects where we can take advantage of abstraction. In the path finding example (section 1.3) we already discussed the advantages of linearly scaling an entire scenario. This is a spatially benign abstraction from absolute size, as the geometric and topological properties of the domain are globally preserved; however, if specific geometric properties such as sizes, distances, and angles change locally, global spatial relations will also be affected. Thus, we may want to have abstractions that only generalize over similar spatial configurations, that is, we may abandon full geometric correctness in favor of a qualitative abstraction that maintains more general spatial characteristics, such as ordering relations and topological relations (Barkowsky et al., 2000).

In addition, we may want to use representations that omit aspects not related to the spatial problem and/or we may want to apply spatially invariant transformations that preserve crucial spatial properties, such as relative distance or relative orientation, and in addition enable new operations. For example, if a geographic problem essentially is a 2-dimensional problem, we can project the 3D spatial environment to a 2D map; by scaling the map from environmental space to the scale of vista space or figural space (Montello, 1993), we permit perceptual and haptic operations on the map that we could not perform on the corresponding entities in geographic space (see section 1.3.2).

3.2 Mild abstraction in geographic maps

In section 1.3 we illustrated for the shortest path problem how we can adapt the representation of a given problem to the problem solving tools available. We argued that for embodied and situated perceiving and acting agents, such as humans, a map representation as a mild abstraction of the spatial environment has advantages over the environment itself, as well as over formal representations in which spatially implicit knowledge is made explicit. In the present section, we will sketch how mild abstraction supports the use of different map types by human users, for solving different types of tasks.

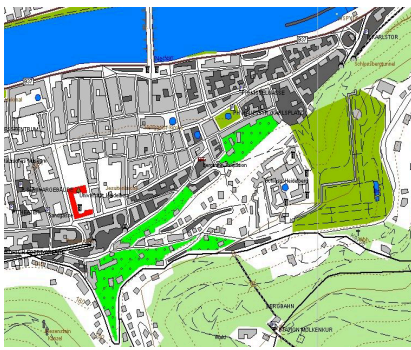
Again, we will maintain the intrinsic properties of spatial structures, where spatial affordances can provide useful shortcuts, while we will abstract from those aspects that benefit from generalization. We will discuss the following familiar types of geographic maps:

- Aerial photograph (abstraction $3D \rightarrow 2D$)
- Topographic map (high-resolution 2D; vertical dimension symbolically represented)
- City map (density of settlements \rightarrow type of urbanization)
- Road map (many spatial features symbolically abstracted (e.g. width of road \rightarrow road type))
- Symbolic sketch map (only topological arrangement spatially represented)

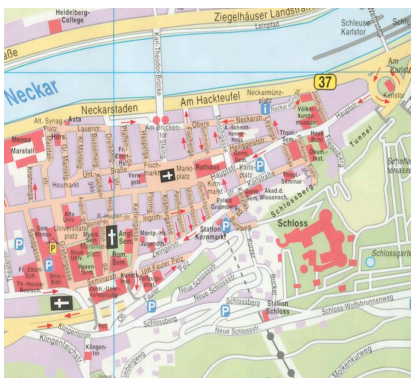
Fig. 7 depicts a section of the city of Heidelberg by each of these map types.



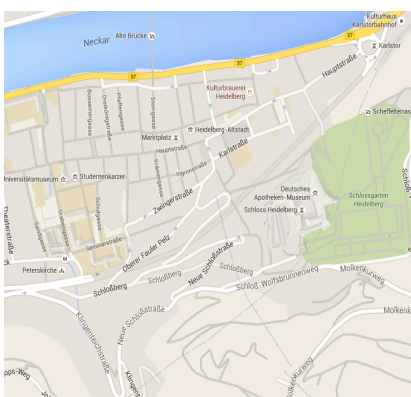
An aerial photograph can be considered the most veridical of the five mappings, as it preserves 2D geometry, as well as other features such as shape, texture, and color to a high degree. Interestingly, this does not imply that for a human in the environment this map is the easiest to match to the environment. Apparently, we can map more easily on a conceptual level than on a severely scaled visual level.



Topographic maps make height information that is lost in the 3D => 2D projection symbolically explicit. Height information is useful for finding hiking routes and scenic views, for example.



Typical city maps make street names and landmarks symbolically explicit by including their names into the map. In addition, churches or other landmarks are symbolically represented. This produces an interface that permits easy matching of entities in the environment with entities in the map.



Typical road maps transform aspects of spatial veridicality into aspects of conceptual identity: Width (and color) of road symbols represents the road type rather than the width of the corresponding road.



Schematic maps of various types emphasize specific aspects of the environment by selectively symbolizing them to support the solution of specific problems. Maps of transportation systems such as subway maps, for example, focus on supporting the transition from one transport line to another. They abstract from other aspects, the representation of which would clutter up the map.

Figure 7. Different levels of mild abstraction for different uses of maps.

In geographic maps, the idea is to maintain as much spatial structure as may be helpful for solving the problem and to add as much symbolic information as may be helpful for dealing with non-spatial aspects of a problem. Spatial information and symbolic information compete for the space of the spatial representation medium. The assessment of the helpfulness can be done on the basis of meta-knowledge through answers to questions such as the following ones:

1. Is the problem given spatially or symbolically?
2. Do we require a spatial or a symbolic solution?
3. Can we treat the problem as an instance of a class of problems for which we know or can find a general approach?
4. Do we know solutions to similar problems?
5. Is the problem at hand peculiar and requires a highly specific approach?
6. Do we need a specific practical (one shot) solution quickly or can we afford to invest time to search for a general solution that may have multiple uses (one-time or long-term optimization)?
7. Is the problem ill-structured or well-structured?

The example of a map may illustrate that, for typical navigation tasks, a reduction from a 3D spatial environment to a 2D projection may be OK, but a reduction to 1D would give up important advantages of spatial integrity. In particular, we lose spatial orientation information when we reduce from 2D to 1D; together with path connectivity, spatial orientation provides important information for selecting suitable path candidates in wayfinding problems. Formally, we can easily construct an information-equivalent representation of all aspects of 3D space by means of 1D relations. In effect, this would correspond to a disintegration of the overall spatial structure; to reconstruct this structure, extensive computation is required.

Thus, in order to decide on a suitable representation for a spatial problem, we should take into account (1) which relations and structures of space we need to *rely on* and (2) which we need to *reason about*. We then can design representations that preserve structures we *rely on implicitly* (like certain distance and orientation relations in a map) and structures that we want to *reason about explicitly*.

3.3 From geographic space to other spatial domains

After illustrating the notion of *mild abstraction* using familiar map representations and familiar wayfinding problems, we will now sketch how we can make use of the concepts described when we want to solve novel spatial problems for which we have to find suitable representations in order to find a solution to the problem. In the following discussion, we will refer to section 2.2, which introduced various components that can be configured for solving spatial problems. As we now have a variety of spatial problem solving approaches at our disposal, meta-knowledge will be helpful to select from the alternatives. We will use spatial puzzles in this section as examples for spatial problems for which successful approaches are not obvious. Spatial puzzles typically are given in the form of a spatial configuration; the goal of the puzzle is to obtain a specific new configuration.

If the objective is to obtain a certain spatial configuration, approach A (ii) (take an action that solves the problem by means of a known effect) is the method of choice. If the necessary knowledge about specific effects of actions is not available, approach A (i) (trial and error) may work if spatial affordances are known and if the number of actions

in space is limited; then the odds may be good enough to find a solution, or the agent can keep track of the steps that have already been attempted.

The approaches A only require the process component 1 (transforming a *spatial* problem into a *spatial* solution). As the problem solving process takes place directly in the spatial medium, transformations into a formal or mental representation of the problem and back are not required. A drawback is that the solution only will fit the given specific problem, as it is not generalized; however, the solution may serve as a source of knowledge for future problems that appear related (for example by means of case-based reasoning (Aamodt & Plaza, 1994)).

If knowledge about the problem is not available, cognitive agents have a choice between approaches A (i) and B (transform spatial problem into a formal problem) – a situation we frequently may be confronted with when we attempt to solve spatial puzzles, such as the snake cube puzzle or Rubik's cube (Fig. 8):



Figure 8. Snake cube puzzle and Rubik's cube: How can we decide how we are going to approach the problem?

Can we solve the problem quickly by trial and error (possibly combined with limited knowledge about effects of local actions) or should we solve the problem via graphical or formal representation and analysis? In addition to process component 2 (transforming a *formal* problem into a *formal* solution), formal analysis requires the process component 3 that transforms the spatial problem into a formal representation of the problem (approach B). This transformation may be hard to find unless we recognize a pattern for which we already have an approach.

In case of graphical or formal analysis we again can distinguish between two cases: B (i) where we know an approach that will solve the corresponding formal problem and B (ii) where we hope to find a solution on the formal level once we have a suitable description of the problem. While B (i) usually makes it easy to find a good problem representation, as the representation is suggested by the known solution to the formal problem, B (ii) can pose severe challenges even for simple-structured problems, as usually there are many ways to represent a given spatial problem and these ways differ in their suitability or ease of solving a given problem. Examples are Geoff Hinton's cube³ and the wine/water mixing problem⁴ (Fig. 9). Both problems are good examples to demonstrate

³ Geoff Hinton's cube: Imagine a cube suspended (by a string) at one of its corners, such that the most distant corner points down vertically. Imagine a vertical axis through these corners. Now turn the cube around this axis in one direction. By how many degrees do you have to turn the cube until for the first time all corners of the cube will coincide with the corners and all edges will coincide with the edges of the cube before rotation? (Hinton, 1979; Freksa et al., 1985).

⁴ Wine/water mixing problem: You have two drinking glasses. One contains wine, the other contains water of equal volume. One spoonful of wine is taken from the wine glass and added to the water in the second glass. Then an equal amount of the water-wine combination is

that the difficult part of solving novel problems is finding a suitable representation (Polya, 1945); once a good representation for the specific problem has been identified, the solution to the problem may become almost trivial⁵.

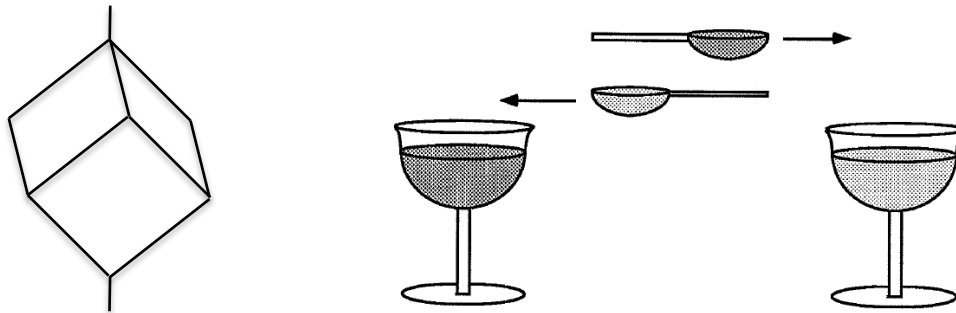


Figure 9 Left: Geoff Hinton's cube.

Right: The wine / water mixing problem.

If the problem solution requires actually carrying out a spatial task as in the spatial puzzles, the process component 4 (transforming formal solution into spatial solution) is also needed; this is usually straightforward if the transformation 3 was performed correctly, as the form of the solution is constrained by the structure of the spatial problem domain. For example, once our formal representation has computed an angle for Hinton's cube, there will be only one axis in the physical cube situation to which this angle may refer, due to the correspondence established in the formalization of the problem.

3.4 Discussion on strategic aspects

The sketched approach to spatial problem solving is particularly economical if we have to solve many spatial problems in the same spatial environment, as the environment needs to be modeled only once for all the problems we have to solve, provided that we represent all the information relevant to these problems and provided that we can represent the environment in such a way that all the problems can be easily solved in that representation. In other words, the spatial problem-solving component 3 (transforming a physically-spatial problem into differently structured representation media) that we did not need for solving the problem directly in the spatial environment may turn out to be a worthwhile investment if (a) we can find a form of representation that can be used for a variety of problems; (b) we need to solve many problems using the same data sets; and (c) we find effective and efficient procedures to solve the problems. Otherwise the search for a suitable problem solving representation may become disproportionally expensive.

transferred from the water glass into the wine glass. Which mixture is purer: the one in the water glass, or the one in the wine glass? (Freksa, 1988).

⁵ Solution to Hinton's cube problem: We can abstract from the vertical dimension and consider a projection of the cube to the horizontal plane. The three edges connecting to the suspending string will divide the 360° space surrounding the axis into three equal sectors of 120°; a turn of 120° will move the cube into complete alignment with its original position. Solution to the wine / water mixing problem: After the transactions, each glass contains the same volume of fluid as before; the volume of wine missing in the wine glass has been replaced by water and vice versa; both volumes are identical. The purity of both fluids therefore is the same.

In today's world of computer power and computerization we are able to use information in so many new ways, that we may be easily led to believe that there will be no limits to solving problems with computers. But we very well know that the computational complexity of some problem classes is so unfavorable, that even simple-structured problems may quickly exceed any computer's capabilities.

If we compare the wayfinding procedure in string maps (section 1.3.2) with wayfinding algorithms we use in computers, we observe that, unlike in computer algorithms, we do not have to consider laws of mathematics to identify a shortest path; consequently, no effort is required to apply such laws to make sure we will compute correct results.

Why do we have to invest in computation in the formalized problem that is not required for solving the original spatial problem? The answer is straightforward: in the spatial environment, all spatial properties relevant to solving spatial problems are *intrinsically* given (Palmer, 1978); that is, they cannot be violated or otherwise overcome. In contrast, in our general computer formalisms we are free to describe many more domains than spatial environments and spatial problems; specifically, we can describe impossible worlds, conflicting situations, and much more. Consequently, we must apply a 'computational straitjacket' to make sure the computer conforms to the laws of physical space when solving spatial problems.

4 Conclusion

A theory of spatial problem solving cannot start and end on the knowledge representation level. It must include the spatial medium, its perception, its representation, the processes operating on the representation, and the spatial actions that result from these processes.

In this chapter we have discussed alternative ways in which cognitive systems can solve spatial problems: (1) by perception and action directly in a spatial medium whose structure reflects the spatial structure of the problem domain; (2) by reasoning in a formal medium in which knowledge about space is explicitly described; or (3) in a combination thereof. As there are several alternatives in which a given problem can be approached, we require an entity that constructs the approach to be taken or that decides on an available selection. This entity will be able to take better decisions the better it understands the problem to be solved.

In classical AI programming, the task of deciding on suitable representations for problem solving usually is carried out by the program designer. His or her creativity constitutes much of the intelligence that is later attributed to the program. If we want to model the creativity of designers of versatile spatial problem solving agents, we should equip spatial problem-solving programs with the knowledge about spatial problem solving that has been described in this chapter.

Fig. 10 depicts three levels of processing that are involved in spatial problem solving:

- (1) The object level is manifested by the substrate that hosts spatial configurations. The substrate effectively integrates (i.e. collocates) the attributes of spatial entities. For example, location, size, orientation, and also molecular structure, color, density, weight, etc. are all integrated at the same place. Spatial configurations of arbitrary entities are the objects of spatial problems and their solutions. Spatial perception, such as vision, haptics, and audition, as well as

spatial actions, such as rotation or motion, operate on the object level. In the end, spatial problems must be solved on the object level. So, the object level is the level on which direct solutions in the spatial medium are performed.

- (2) The knowledge level *describes* the object level. It makes knowledge about relevant properties of the substrate and the spatial configurations explicit, typically in the form of statements about facts and relations. Different aspects of the object level (e.g. location, size, orientation) are described individually in a linear fashion, i.e. in some formalism or as text. This means that the spatial structure of the object level is disintegrated into the various aspects that are made explicit. Explicit knowledge is the basis for argumentation and reasoning. Thus, we may have inference rules and calculi on the knowledge level that enable us to derive new facts and relations about the object level. So, the knowledge level is the level on which indirect solutions in abstract media are computed.
- (3) The strategy level contains meta-knowledge about actions and affordances on the object level, about rules and processes on the knowledge level, as well as about their effects. The strategy level controls perception, action, problem and knowledge transformation, as well as reasoning for solving the spatial problem at hand. Thus, it decides about the distribution of tasks between the object level and the knowledge level. By maintaining certain spatial structures and ‘aspectualizing’ others (Bertel et al. 2004), the strategy level can make use of the advantages of mild abstraction.

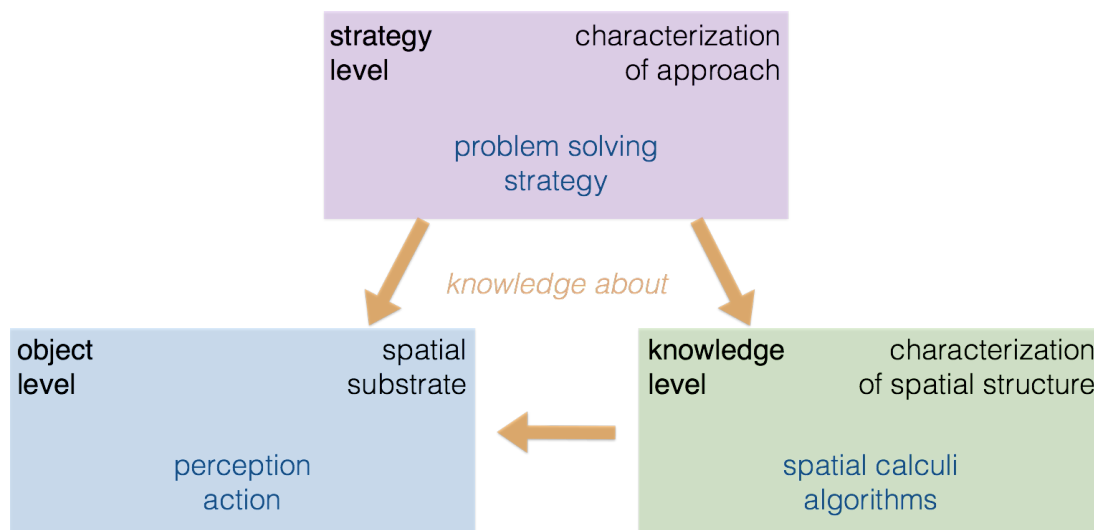


Figure 10: Levels of cognitive processing in spatial problem solving.

We suggest that physical spatial structures play a special role in cognitive processing, as they are capable of integrating multiple aspects of a domain at the same place and of manipulating them simultaneously. We demonstrate how spatial structures can be exploited in spatial problem solving and how they can be combined with more abstract knowledge processing approaches.

At this time, we do not have computer architectures that replicate integrated spatial structures as described in this chapter. We therefore propose to investigate the *strong spatial cognition* paradigm in physical space by means of embodied and situated cognitive agents, such as people and robots, to better understand how to make use of integrated spatial structures in cognitive processing.

Acknowledgments

We acknowledge generous support by the German Research Foundation (DFG SFB/TR 8 projects R1 and R3; DFG project SOCIAL—FR 806/15) and by the Central Research Development Fund of the University of Bremen (project CogQDA). We thank Holger Schultheis, Ahmed Loai Ali, and the editors of this volume for critical and constructive comments.

References

- Aamodt, A., & Plaza, E. (1994). Case-based reasoning: Foundational issues, methodological variations, and system approaches, *Artificial Intelligence Communications*, 7 (1), 39–52.
- Barkowsky, T., Latecki, L. J., & Richter, K.-F. (2000). Schematizing maps: Simplification of geographic shape by discrete curve evolution. In C. Freksa, W. Brauer, C. Habel, & K. F. Wender (Eds.), *Spatial Cognition II – Integrating Abstract Theories, Empirical Studies, Formal Models, and Practical Applications* (41–53). Berlin: Springer.
- Berendt, B., Barkowsky, T., Freksa, C., & Kelter, S. (1998). Spatial representation with aspect maps. In C. Freksa, C. Habel, & K. F. Wender (Eds.), *Spatial cognition – an interdisciplinary approach to representing and processing spatial knowledge* (313–336). Berlin: Springer.
- Bertel, S. (2010). *Spatial Structures and Visual Attention in Diagrammatic Reasoning*. Lengerich: Pabst Science Publishers.
- Bertel, S., Freksa, C., & Vrachliotis, G. (2004). Aspectualize and conquer. In J. S. Gero, B. Tversky, & T. Knight (Eds.), *Visual and Spatial Reasoning in Design III* (255–279). Sydney: Key Centre of Design Computing and Cognition, University of Sydney.
- Bobrow, D. G. (1975). Dimensions of representation. In D. G. Bobrow, & A. Collins (Eds.), *Representation and Understanding* (1–34). New York: Academic Press.
- Card, S. K., Mackinlay J. D., & Shneiderman, B. (1999). *Readings in Information Visualization: Using Vision to Think*. San Francisco: Morgan Kaufmann.
- Davis, E. (2013). Qualitative spatial reasoning in interpreting text and narrative. *Spatial Cognition and Computation*, 13 (4), 264–294.
- Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, 1, 269–271.
- Dirlich, G., Freksa, C., & Furbach, U. (1983). A central problem in representing human knowledge in artificial systems: The transformation of intrinsic into extrinsic representations. *Proc. 5th Cognitive Science Conference*. Rochester: Cognitive Science Society.
- Dreyfus, H., & Haugeland, J. (1974). The computer as a mistaken model of the mind. In S. C. Brown (Ed.), *Philosophy of Psychology* (247–258). London: Palgrave Macmillan.
- Euclid (1956). *The Thirteen Books of Euclid's Elements*. (T. L. Heath, Trans.). New York, NY: Dover. (Original work published 300 BC).
- Fikes, R. E., & Nilsson N. J. (1971). STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2, 189–208.
- Freksa, C. (1988). Intrinsische vs. extrinsische Repräsentation zum Aufgabenlösen oder die Verwandlung von Wasser in Wein. In G. Heyer, J. Krems, & G. Görz (Eds.), *Wissensarten und ihre Darstellung* (155–165). Berlin: Springer.
- Freksa, C. (1991). Qualitative spatial reasoning. In D. M. Mark, & A. U. Frank (Eds.), *Cognitive and Linguistic Aspects of Geographic Space* (361–372). Dordrecht: Kluwer.
- Freksa, C. (1997). Spatial and temporal structures in cognitive processes. In C. Freksa, M. Jantzen, & R. Valk (Eds.), *Foundations of Computer Science. Potential – Theory – Cognition* (379–387). Berlin: Springer.
- Freksa, C. (1999). Spatial aspects of task-specific wayfinding maps: A representation-theoretic perspective. In J. S. Gero, & B. Tversky (Eds.), *Visual and Spatial Reasoning in Design* (15–32). Sydney: Key Centre of Design Computing and Cognition, University of Sydney.

- Freksa, C. (2015). Strong spatial cognition. In S. I. Fabrikant, M. Raubal, M. Bertolotto, C. Davies, & S. Bell (Eds.), *Spatial Information Theory* (65–86). Heidelberg: Springer.
- Freksa, C., Furbach, U., & G. Dirlich (1985). Cognition and representation. In J. Laubsch (Ed.), *GWAI-84. 8th German Workshop on Artificial Intelligence* (119–144). Berlin: Springer.
- Freksa, C., Moratz, R., & Barkowsky, T. (2000). Robot navigation with schematic maps. In E. Pagello, F. Groen, T. Arai, R. Dillmann, & A. Stentz (Eds.), *Intelligent Autonomous Systems, 6*, 809–816. Amsterdam: IOS Press.
- Freksa, C., Oltejeanu, A. M., Ali, A. L., Barkowsky, T., van de Ven, J., Dylla, F., & Falomir, Z. (2016). Towards spatial reasoning with strings and pins. *Advances in Cognitive Systems 4 Poster Collection #22*, 1-15.
- Furbach, U., Furbach, F., & Freksa, C. (2016). Relating strong spatial cognition to symbolic problem solving – An example. *Proc. 2nd Workshop on Bridging the Gap Between Human and Automated Reasoning*, IJCAI, New York, 9 July 2016, arXiv:1606.04397 [cs.AI]
- Gärdenfors, P. (2000). *Conceptual Spaces*. Cambridge, MA: MIT Press.
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. New Jersey: Lawrence Erlbaum Assoc.
- Hinton, G. (1979). Some demonstrations of the effects of structural descriptions in mental imagery. *Cognitive Science*, 3, 231–250.
- Johnson, S. P. (2009). Developmental origins of object perception. In A. Woodward, & A. Needham (Eds.), *Learning and the Infant Mind* (47–65). New York: Oxford University Press.
- Keen, R. (2003). Representation of objects and events: Why do infants look so smart and toddlers look so dumb? *Current directions in Psychological Science*, 12, 79–83.
- Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11, 65–100.
- MacEachren, A. M. (1995). *How Maps Work: Representation, Visualization, and Design*. New York, NY: Guilford Press.
- Mark, D. M., Freksa, C., Hirtle, S. C., Lloyd, R., & Tversky, B. (1999). Cognitive models of geographic space. *Int. J. of Geographical Information Science*, 13, 8, 747–774.
- Marr, D. (1982). *Vision*. Cambridge, MA: MIT Press.
- McCarthy, J., & Hayes, P. J. (1969). Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence*, 4, 463–502.
- Monmonier, M. S. (1996). *How to Lie with Maps*, 2nd ed. Chicago IL: University of Chicago Press.
- Montello, D. (1993). Scale and multiple psychologies of space. In A. U. Frank, & I. Campari (Eds.), *Spatial Information Theory: A Theoretical Basis for GIS* (312–321). Berlin: Springer.
- Needham, A. (2009). Learning in infants’ object perception, object-directed action, and tool use. In A. Woodward, & A. Needham (Eds.), *Learning and the Infant mind* (208–226). New York: Oxford University Press.
- Oltejeanu, A. M. (2016). *A Cognitive Systems Framework for Creative Problem Solving*. Doctoral dissertation. Bremen: University of Bremen.
- Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch, & B. B. Lloyd (Eds.), *Cognition and Categorization*, 259–303. Hillsdale: Lawrence Erlbaum.
- Polya, G. (1945). *How to Solve it*. Princeton: Princeton University Press.
- Robinson, A. H., Morrison, J. L., Muehrcke, P. C., Kimerling, A. J., & Guptill, S. C. (1995). *Elements of Cartography*, 6th edition. New York: Wiley.
- Russell, S., & Norvig, P. (1995). *Artificial Intelligence -- A Modern Approach*. Upper Saddle River: Prentice Hall.
- Scaife, M., Rogers, Y. (1996). External cognition: How do graphical representations work? *Int. J. Human-Computer Studies*, 45, 185–213.
- Sloman, A. (1971). Interactions between philosophy and artificial intelligence: The role of intuition and non-logical reasoning in intelligence, *Artificial Intelligence*, 2, 209–225.

- Sloman, A. (1985). Why we need many knowledge representation formalisms. In M. A. Bramer (Ed.), *Research and Development in Expert Systems*, 163–183. Cambridge: Cambridge University Press.
- Tversky, B. (1981). Distortions in memory for maps. *Cognitive Psychology*, 13, 407–433.
- Tversky, B. (1992). Distortions in cognitive maps. *Geoforum*, 23, 131–138.
- Tversky, B. (1993). Cognitive maps, cognitive collages, and spatial mental models. In A. U. Frank & I. Campari (Eds.) *Spatial Information Theory: A Theoretical Basis for GIS*, 14–24, Berlin: Springer.
- Tversky, B. (2000). Some ways that maps and diagrams communicate. In C. Freksa, W. Brauer, C. Habel, & K. F. Wender (Eds.), *Spatial Cognition II – Integrating Abstract Theories, Empirical Studies, Formal Models, and Practical Applications*, 72–79, Berlin: Springer.
- Tversky, B., & Lee, P. (1999). Pictorial and verbal tools for conveying routes. In C. Freksa, & D. Mark (Eds.), *Spatial Information Theory*, 51–64, Berlin: Springer.
- Yan, J., Forbus, K., & Gentner, D. (2003). A theory of rerepresentation in analogical matching. *Proc. Twenty-fifth Annual Meeting of the Cognitive Science Society*, New York: Psychology Press.