D.M. Mark & A.U. Frank (eds.), Cognitive and Linguistic Aspects of Geographic Space, 361-372.

QUALITATIVE SPATIAL REASONING

CHRISTIAN FREKSA*
Institut für Informatik
Technische Universität München
Arcisstr. 21
8000 München 2
Germany

ABSTRACT. Physical space has unique properties which form the basis of fundamental capabilities of cognitive systems. This paper explores some cognitive aspects of perception and knowledge representation and explains why spatial knowledge is of particular interest for cognitive science. It is suggested that 'spatial inference engines' provide the basis for rather general cognitive capabilities inside and outside the spatial domain. The role of abstraction in spatial reasoning and the advantages of qualitative spatial knowledge over quantitative knowledge are discussed. The usefulness of spatial representations with a low degree of abstraction is shown. An example from vision (the aquarium domain) is used to illustrate in which ways knowledge about space may be uncertain or incomplete. Parallels are drawn between the spatial and the temporal domains. A concrete approach for the representation of qualitative spatial knowledge on the basis of 'conceptual neighborhood' is suggested and some potential application areas are mentioned.

1. What is special about spatial knowledge?

Our knowledge about physical space differs from all other knowledge in a very significant way: we can perceive space directly through various channels conveying distinct modalities. Unlike in the case of other perceivable domains, spatial knowledge obtained through one channel can be verified or refuted by perception through the other channels. As a consequence, we are disproportionally confident about what we know about space: we take it *for real.*

More specifically, we perceive spatial dimensions visually, tactily, acoustically, and even by smell and by temperature sensation. Some animals exploit perception of electrical or magnetic fields for spatial orientation. Within the visual sense, we obtain clues about space through size, brightness, hue, saturation, and texture information; within the acoustic sense, we obtain clues about space through loudness, frequency, and signal/noise ratio. In addition to perceiving space multimodally, we can modify spatial situations by moving objects, we can expose objects to the perceivable influence of physical forces, and we can modify our spatial perception by moving ourselves. In this way, we experience the laws of space as we can experience no other dimension.

In contrast, the perception of other dimensions like color, visual brightness, sound, smell, temperature, etc. relies on single modalities and we cannot directly modify them. This makes the

---

* Research supported by Deutsche Forschungsgemeinschaft under grants Br 609/4-2, Fr 806/1-1, and by Siemens AG.

knowledge about these dimensions much more succeptible to perceptional illusion or uncertainty. I will argue that physical space plays a central role in cognition 1) as the domain in which physical events take place, and 2) as reference domain for the interpretation of non-spatial concepts.

The special cognitive reality of space described above makes the spatial domain particularly suitable as a medium for conveying knowledge, since its properties are universal to different cognitive systems. Thus, the spatial domain can be used particularly well as the source domain for metaphors with a non-perceivable or abstract target domain. In this way, the properties of physical space can be used as a vehicle for conveying non-spatial concepts, provided there exists a mapping from the non-spatial to the spatial domain. Lakoff & Johnson [1980] and Lakoff [1987] present numerous examples for spatial metaphors in natural language.

I propose that our knowledge about the organization of space serves as a ‘cognitive interface’ between abstract or non-perceptual knowledge and the ‘real world’. In other words, we may interpret non-spatial concepts by mentally transforming them into spatial concepts (i.e., *understanding* them in terms of spatial concepts), carrying out mental operations in this ‘visualizable’ and ‘graspable’ domain and transforming the result into the original domain. In this way, ‘spatial inference engines’ may have a much more general function than the term suggests: rather than generalizing by forming a common abstraction for various domains we generalize by forming suitable analogies to a well-understood concrete domain.

2. Spatial inference engines

If there is a large class of problems which can be solved well under the specific constraints of the spatial domain, we may want to develop a specialized representation system with an ‘inference engine’ which is capable of dealing with knowledge about space particularly well. The advantage of a special purpose approach to spatial reasoning is that certain constraints which always hold in the spatial domain do not have to be modeled and verified in each situation anew.

The spatial domain is severely constrained in certain ways in which abstract ideas are not. For example, distinct solid objects never occupy the same physical space and a single physical object always occupies a single contiguous space. The arrangements of objects and locations given by their neighborhood relations are very important features of spatial situations which are relevant in virtually all spatial problem solving tasks. This has to do with the fact that physical objects consist of connected parts (‘neighbors’) and with the fact that movement in space is only possible between neighboring locations. Therefore we do not want to check in a general representation explicitly if these constraints hold; the constraints should be ‘built-in’ in the reasoning system as they are built-in in the spatial domain.

For spatial problem solving with a general-purpose theorem prover, for example, such constraints must be explicitly formulated (compare Reiter & Mackworth [1989]). This is due to the extreme degree of abstraction from the original problem which is possible and necessary in a general-purpose representation system. In reasoning about concrete domains like physical space, however, such a high level of abstraction is not desirable since additional effort is required to nullify excessive abstraction as in the example given above. Instead, we are aiming at a representation level which only abstracts from unnecessary details but maintains the important aspects of the domain.

This approach is not new, of course: constructive geometry exploits spatial properties of drawing paper as an analogy to properties in the abstract model domain; cartography utilizes spatial properties of maps as an analogy to spatial properties in the depicted real world; fluid mechanics makes use of self-similarity of space to form an analogy between small scale space and larger scale space when wind tunnel experiments on small scale models are used to predict the behavior of larger scale objects. In addition, neurobiological research has revealed that various areas of the visual cortex are arranged in such a way that certain spatial properties of the external world are preserved in the cortical structures. The crucial question for artificial intelligence in this context is how we can implement a spatial reasoning system which makes direct use of space and its intrinsic properties.

3. The role of abstraction for spatial reasoning

If we want to maintain properties of space for spatial reasoning, why not solving spatial problems in the real world directly? In fact, there are situations in which this is probably the best approach. For example, for identifying an object in an unusual position or orientation it may be easier to adjust position or orientation than to formalize the object and mentally match the representation against the representation of a template. Also, whenever there is insufficient knowledge for formalization, the *world-based* approach is most promising: rather than using abstract knowledge about the world and its properties, ‘real world tasks’ can be solved directly under the influence of the world’s real properties. Examples are: finding the right key for a lock; testing a vehicle under realistic conditions; diagnosing satellite malfunctions by simulating the malfunctioning process with the aid of a duplicate satellite.

However, in many spatial situations it is too expensive to determine solutions in the real world directly. This is the case when large or heavy objects or large distances are involved. In such cases, we would like to abstract from actual size, weight, or distance; but we have no need to abstract from fundamental spatial properties like spatial arrangement. The usual approach in this situation is to formalize knowledge about the abstract problem and to solve it on a formula manipulation device, a computer. The reason for representing *knowledge* about a domain rather than representing the domain itself is that the world is accessible by formal means only through knowledge. When we are using computers we are bound to using formal approaches to modelling reasoning processes. At this point, however, we do not have formal languages for computers which are suitable for the level of abstraction we described above.

Formal languages like predicate logic initially abstract from all inherent properties of space. Specifying the spatial domain in such an abstract language means constraining the language in order to make the semantics consistent with the target domain ‘space’. We are aiming at a language for representing spatial knowledge which leaves us only the expressive freedom we can utilize in the spatial domain (‘spatial inference engine’ level in Figure 1). In other words, we do not want to specify the laws of space but we want to take them for granted.

Abstraction is very useful in its own right: one main attraction of abstraction is that it liberates representations from insignificant details and focusses on the significant distinctions. This advantage, however, gets partially lost if we go overboard by using formalisms which liberate the representation from significant constraints as well.

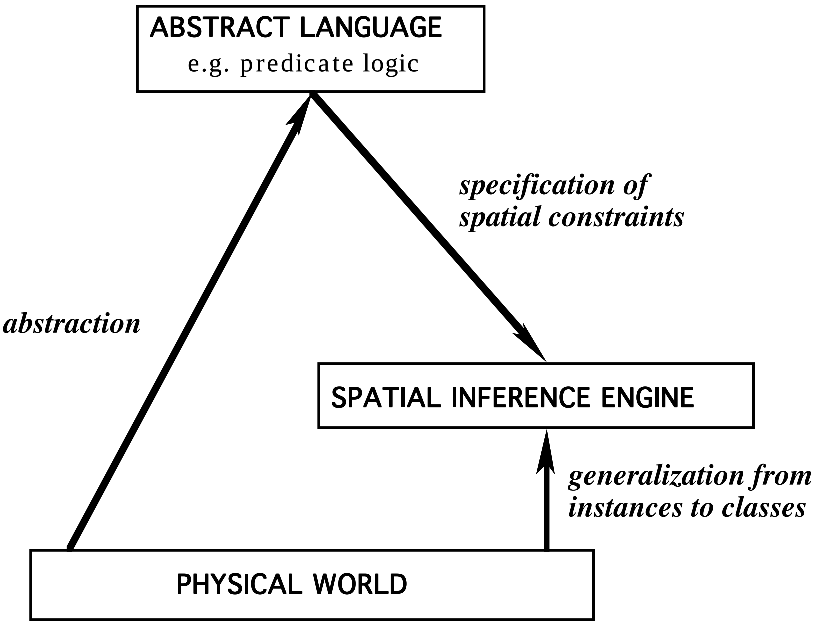

Figure 1: Spatial inference engines can be constructed by constraining very abstract formalization languages or by generalizing over the physical world.

3.1. WHICH PROPERTIES OF REAL SPACE DO WE WANT TO ABSTRACT FROM?

Ideally, we want to represent only knowledge which we will need for solving the tasks which we have to solve and abstract from all other knowledge. Since we typically collect knowledge before we know the specific tasks for which we will use it we can not be so restrictive; but depending on the classes of tasks we want to solve we are usually able to exclude a range of items which we will never consider.

For example, in the domain of spatial knowledge it makes a difference if we represent knowledge in order to construct a precise metric map of some area or if we represent knowledge in order to describe a location to be identified in the real world as in the case of wayfinding. In the first case, we require quantitative knowledge which enables us without reference to the real world to predict precisely at which (absolute) location we will encounter which object. In the second case, we only require enough knowledge to distinguish quantities which are given in the real world and which are available for reference, so a limited amount of qualitative knowledge may suffice.

We are particularly interested in the second case, i.e. the class of identification tasks, since many spatial reasoning problems can be reduced to this class of tasks. In fact, all reconstruction tasks can be structured hierarchically into identification subtasks. We can learn about useful and cognitively adequate abstractions by considering established methods of dealing with spatial knowledge. In the domains of object-, scene-, or route description and identification, for example, we can refer to the long tradition of map production and look which aspects of the spatial domain have been found to be significant in good maps.

For description and identification tasks our culture typically uses linguistic and/or pictorial means to convey knowledge, i.e., we use a natural language description or a depiction of the situation, or preferably a combination of the two. In both approaches, part of the knowledge to be conveyed is used to establish the context of the situation and part of the knowledge actually establishes the target object within that context. A main advantage of combining language and depiction is to reduce the description itself to the identification task level.

3.2. WHICH PROPERTIES OF REAL SPACE DO WE WANT TO MAINTAIN?

For spatial reasoning tasks, the following conceptual properties of spatial configurations appear to be universally useful; we will argue that they should be preserved by the representational structure.

1) Universal domain properties which do not depend on specific situations, for example
    - each object exists exactly once
    - each location coincides with at most one object

   (uniqueness constraints)
2) Properties related to the physics of space (arrangement information):
    - movement in space is possible only between neighboring locations
      (this property should hold on any level of representation)

   (topology)
3) Properties related to the neighborhood of spatial relations (compare Freksa [1990]). For example, the spatial relations described by ‘A is left of B’ and ‘A touches B on the left’ are spatial-conceptual neighbors, since there is a direct transition from one relation to the other when the spatial arrangement is modified; in contrast ‘A is right of B’ is not a spatial-conceptual neighbor, since other relations have to be traversed in the transition.

   (conceptual structure)

The first two properties are the same properties which constitute the usual semantics of sketch maps of spatial environments (compare Reiter & Mackworth [1989]): each object is depicted only once; different objects are positioned at different locations, and the objects are arranged according to their arrangement in the real world. Other properties of sketch maps are not considered significant, e.g. the thickness of the lines, the exact lengths of the lines, the precise orientation of the lines, the completeness of the representation. The third property (conceptual structure) becomes particularly important for reasoning under uncertainty or in absense of complete knowledge (compare Freksa [1990]).

3.3. QUALITATIVE KNOWLEDGE

Much of the knowledge about time and space is qualitative in nature. Specifically, this is true for visual knowledge about space. Although the retinal image of a visual object is a fully instantiated quantitative image in the sense that specific locations on the retina are stimulated by light of a specific spectrum of wavelengths and intensity, the knowledge about a retinal image that can be retrieved from memory is qualitative. Absolute locations, wavelengths, and intensities cannot be retrieved from memory, in general; rather, comparisons between such features can be performed. Thus, from the representation of the image, the original retinal image cannot be reconstructed in full; merely certain qualitative relations among features within the image or between image features and memory features can be recovered.

Qualitative representations of the kind characterized above share a variety of properties with ‘mental images’ (see, for example, Block [1981]) which people report about when they describe from memory what they have seen or when they attempt to answer questions on the basis of visual memories:

- depending on familiarity or previous knowledge they will be able to retrieve fewer or more details;
- in the absence of related features, descriptions will be rather vague;
- the accessible information is not sufficient for reconstructing the original image accurately to scale; rather certain qualitative relationships can be reconstructed;
- if enough such qualitative relationships can be recalled, a fairly accurate reconstruction of the original image is possible due to the interaction of constraints;
- quantitative knowledge can be represented by ‘anchored’ qualitative knowledge (≡ measuring);
- qualitative knowledge is robust under transformations.

I would like to suggest that qualitative knowledge of this kind is exactly the type of knowledge we want to represent for spatial reasoning and what we need for solving identification tasks. Most qualitative relations act as constraints which – individually viewed – leave quite a bit of freedom as to the actual quantitative values possible to satisfy the constraints; some qualitative constraints are as informative as quantitative ones, however. For example, the assertion ‘a is smaller than b’ describes a qualitative constraint which – in a given domain – may be satisfied by a very large number of quantities a, given a certain value for b. In contrast, the qualitative relation described by ‘a equals b’ constrains the quantity of a to the single value of b. Similarly, the assertion ‘k is perpendicular to l’ constrains the orientation of a directed path in a 2-dimensional domain to exactly two possible values.

The expressive power of qualitative constraints results from their interaction. For example, the assertions ‘a is smaller than b’ and ‘b is smaller than c’ constrain the value of b to the values of the interval]a, c[. In the domain of real values there is still an infinite number of quantities in this interval; however, if we further constrain the assertions to refer to a domain of discrete entities, the qualitative constraints may have the power of selecting small sets of quantities without directly addressing their value.

For solving identification tasks we are not interested in the theoretical specificity of descriptions as obtained by considering the potential set of entities referred to by a descriptor; rather, the *actual* set of entities and values that exist in the specific domain is of interest. Qualitative descriptions are very powerful in this situation: when there are many items to be distinguished there are many relations that can be established for distinguishing them; when there are few, only few are needed for their distinction.

A great advantage of qualitative representations for cognitive systems is the fact that they are independent of specific values and granularities of representation; depending on the situation context and on the granularity of the available knowledge they correspond to more specific or more general entities. In this way, qualitative representations allow for a top-down approach to characterizing situations, in comparison to bottom-up approaches suggested by quantitative representations.

For example, suppose we want to describe objects in terms of their sizes. If we represent sizes quantitatively, we express their values in terms of some absolute unit. Depending on the granularity of the unit chosen, we can distinguish fewer or more classes of sizes; thus, we must know beforehand which size differences we may want to distinguish. If we represent sizes qualitatively, e.g. in terms of the relations 'smaller', 'equal', 'larger', we do not predefine the size differences we will be able to distinguish; depending on the specific needs, we can compose coarser or finer descriptions by cascading relations, i.e., by making fewer or more comparisons of the quality involved. In the absense of a basic reference unit, 'granularity' becomes a relative concept; coarseness of a representation depends on the context of concepts involved.

Qualitative representations structure a domain according to conceptual distinctions corresponding to the way the domain is viewed rather than according to physical quantities the domain is measured by. In other words, qualitative representations do not structure domains homogeneously (i.e., with uniform granularity of physical entities) as quantitative representations do; rather, they focus on the boundaries of concepts: the representation may be viewed as having low resolution for different values corresponding to the same quality and high resolution near the concept boundaries. This is true at least for crisp concepts, i.e., concepts with sharp boundaries. Thus, qualitative representations may be viewed as regions from the viewpoint of quantitative representations and as nodes from a conceptual viewpoint.

4. The aquarium metaphor for spatial concepts

I will illustrate some properties of perceptual knowledge which we typically encounter when we represent knowledge about the *real world* by introducing the 'aquarium metaphor' for communication about spatial perceptions (compare Freksa [1980]): Two observers A and B look at an aquarium and get excited about the beautiful fishes they see (Figure 2). They do not know the names of the fishes but they want to communicate about particular fish individuals to one another.

The problem observer A are confronted with is to describe the fish in such a way that observer B is able to identify the fish correctly. The situation can be characterized as follows:

– A and B cannot quantitatively locate the fish (because the glass prevents them from putting a ruler inside the aquarium, paralax prevents them from aligning a ruler outside the aquarium, and the movement of the fish prevents them from watching it long enough to precisely determine its position otherwise.

\+ However, the observers are not completely uncertain about the position of the fish. Specifically, they have knowledge about its position *relative* to other objects (qualitative knowledge); furthermore, these relative positions do not change arbitrarily when the fishes move; if the relation between two fishes changes at all, a neighboring relation will replace the previous relation (conceptual neighborhood of relations).

– There are various resource limitations: the amounts of time for observation and for object identification are limited, perceptual resolution is limited (coarse knowledge), and the perceivable features are limited (incomplete knowledge).

\+ As a consequence, the complexity of the description and identification tasks is limited.

– The movement and muddiness of the water may prevent the observers from clearly recognizing the boundaries of objects and relations (fuzzy knowledge).

+ They can use fuzzy natural language descriptors which combine conceptually neighboring situations corresponding to their observations (adaptivity of natural language descriptors).
– The two observers derive their knowledge about the aquarium world from different perceptions, since they are located at different observer positions; as a consequence, some objects and features may be partially or completely hidden from one observer but not from the other and the observers may observe different spatial relations for the same spatial situation (subjective knowledge).
+ The knowledge of the two observers is strongly related: gradual changes in the observer position cause at most stepwise changes between neighboring spatial relations; one observer can easily model how the observation of the other observer may deviate from the own observation (conceptual neighborhood of relations).
+ The two observers are not only related by the spatial closeness of their positions but also by the common situation context to which they refer. Thus, the object description and identification task is strongly simplified since identification must be possible only in relation to this context (context-aided communication).

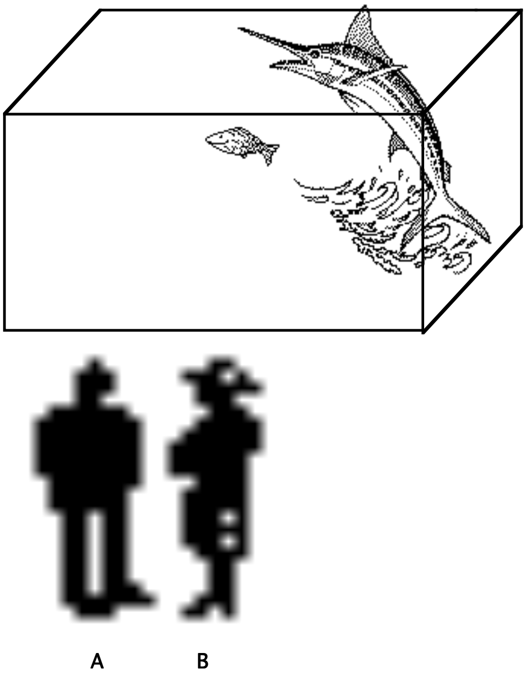

Figure 2: Observer A communicates with observer B about a particular fish in the aquarium.

In other words, the knowledge about the aquarium world which the two observers can use for object description and identification is incomplete, imprecise, and subjective. Nevertheless, the observers will be able to convey to each other which objects in the aquarium world they are focussing at: they can describe and locate the objects by qualitative means, i.e., they can observe a number of relations between features of the target object and some other objects in the aquarium and convey their knowledge in terms of these relations; they can specify a few features which may suffice to distinguish the target object from the other objects; they will be able to interpret descriptions 'with a grain of salt', i.e., if the description does not completely fit their own conception of the situation they will be able to relax the interpretation in such a way that crucial distinctions are still maintained and less crucial distinctions are tolerated.

The aquarium metaphor emphasizes a few issues which are present in all spatial perception, representation, and identification situations:

1) perceptual knowledge is necessarily limited with regard to resolution, features, completeness, certainty;
2) there is a more or less well-defined context;
3) perceptions are finite;
4) the neighborhood of objects and the conceptual neighborhood of relations between objects provide very useful information for spatial reasoning

In addition to these perception-related issues, the aquarium metaphor is useful for studying language-related issues of physical space. For example, which properties of space are preserved implicitly, which ones are represented explicitly, and which ones are not represented at all? What are the common features and what are the differences between a linguistic description and, say, a sketch map of the same scene?

Note that the aquarium metaphor is not universal: for example, it does not fit an abstract view of space in which the spatial relations 'right of' and 'left of' may be considered closely related, because they are symmetric. As a perception-oriented construct, the aquarium metaphor is very close to the physics of space, but not to arbitrary abstractions thereof.

In our research, we focus on spatial representation, description, and identification tasks which share the properties discussed in the framework of the aquarium metaphor.

5. The one-dimensional case

James Allen [1983] proposed an approach to representing temporal knowledge which is based on qualitative information about time. More specifically, Allen represents relative positions between events represented as intervals. In the linear case of time, any two events have exactly one out of 13 possible mutually exclusive qualitatively distinct relationships to each other (the orientation of the objects being ignored). In Figure 3, these relationships are shown for the one-dimensional spatial case: two spatial objects, here a black and a grey fish, are shown in 13 qualitatively different relative positions with regard to the horizontal dimension.

The thirteen relations are denoted by the labels *before* (<), *after* (>), *during* (d), *contains* (di), *overlaps* (o), *overlapped-by* (oi), *meets* (m), *met-by* (mi), *starts* (s), *started-by* (si), *finishes* (f), *finished-by* (fi), and *equals* (=). In the figure, the 13 cases are arranged according to 'conceptual neighborhood', i.e., relations between which direct transitions are possible when a fish moves or changes its size are located next to each other.

In Allen's scheme, reasoning is done by computing the set of possible relationships between two entities whose relationships to a third entity must be known. In Freksa [1990], Allen's approach has been generalized to allow for reasoning with incomplete or coarse knowledge about temporal relations. 'Conceptual neighborhood' of relations is the key concept for this generalization. It allows for inferring neighborhoods of relations between objects about which neighborhoods of relations to some other object are known. Interactions of such sets of relations may lead to rather precise inferences. The inferences that can be drawn are not in all cases completely precise – some ambiguity as to which of a neighborhood of relations holds may remain. No quantities are involved in the representation and reasoning can be carried at any level of granularity. Arbitrary levels of granularity can be mixed. The resulting qualitative conclusions are always correct, although – by necessity – they may be rather coarse if translated into quantities. The approach fulfills the requirements for solving identification tasks as discussed in connection with the aquarium metaphor in the previous section. Our aim is now to adapt this approach to higher-dimensional physical space.

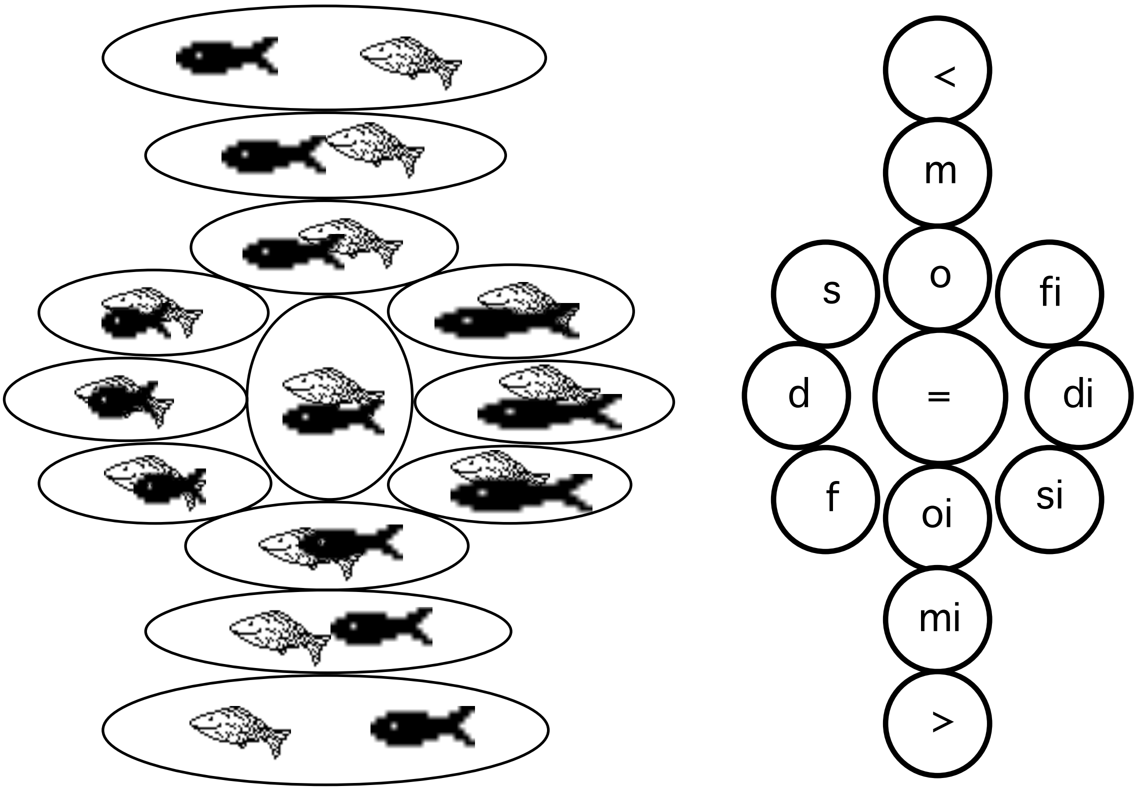

Figure 3: Left: One-dimensional spatial relations between two objects (fishes) arranged according to their conceptual neighborhood. Right: The corresponding labels arranged accordingly.

---

6. Higher-dimensional space

The spatial domain is considerably more complex than the temporal domain for which the qualitative reasoning scheme originally was developed: the everyday concept of time is one-dimensional and uni-directional; if we give up uni-directionality but add orientation of objects in one dimension, we obtain 26 relations between two objects: each relation in Figure 3 obtains a counterpart where the two fishes face the opposite direction. When moving to 2- or 3-dimensional spaces, in theory more qualitative relations between two objects can be distinguished (compare Güsgen [1989]); however, it is important to determine which of those relations are meaningful in the specific spatial domain under consideration and at which stage of reasoning the distinctions between certain relations become significant. The combination of conceptually neighboring relations at certain stages of reasoning may simplify the representation and/or the reasoning processes considerably (compare Freksa [1990]).

More specifically, we must determine

1) which type of objects do we want to represent
   a) only solid objects or objects which can penetrate each other?
   b) only convex objects or also concave objects?
   c) blocks, ellipsoids, sectors, or arbitrary shapes?
   d) the objects themselves or views of the objects (e.g. projections)?
2) which type of reference systems we need
   a) deictic, intrinsic (with reference to single objects), extrinsic (with reference to relations between objects)?
   b) are objects viewed as entities without orientation or as entities with orientation; what is the reference frame for the orientation?
3) which type of knowledge organization we need
   a) flat or hierarchical organization?
   b) level-independent or level-specific organization?

Depending on the specific decisions taken, different structures are advantageous and different complexities result. Hernández [1990] presents a representation scheme for 2-D projections of 3-D convex objects (“abstract maps”) with intrinsic, extrinsic, or deictic reference system; the knowledge is hierarchically structured with a level-specific organization.

7. Application areas for qualitative spatial reasoning

The paper started off with a motivation for developing spatial inference engines as rather universal vehicles for cognition. In the subsequent chapters we discussed issues specifically relevant to the spatial domain. Once one’s attention is focussed on spatial reasoning it appears to be omnipresent. In our research at the Technische Universität München we are presently concerned with reasoning about *real* (i.e. non-metaphorical) space, specifically with description of spatial scenes and with interpretation of scene descriptions with reference to a representation of the scene. This type of reasoning is particularly relevant to route description and to navigation. Spatial planning is another area for which qualitative spatial reasoning is appropriate. Qualitative spatial reasoning also is useful for diagnosis of malfunctions in electrical circuits or in neurology and for Geographic Information Systems (GIS).

Acknowledgements

I acknowledge numerous discussions with members and students of the Artificial Intelligence/ Cognition research group at the Technische Universität München and with participants of the Advanced Study Institute at Las Navas del Marqués.

References

Allen, J., 1983. Maintaining knowledge about temporal intervals. CACM 26, 11, 832-843.

Block, N., ed., 1981. Imagery. MIT Press. Cambridge, Massachusetts.

Freksa, C., 1980. Communication about visual patterns by means of fuzzy characterizations. XXIInd Intern. Congress of Psychology, Leipzig.

Freksa, C., 1990. Temporal reasoning based on semi-intervals. ICSI TR-90-016, International Computer Science Institute, Berkeley.

Freksa, C., Habel, C., 1990. Warum interessiert sich die Kognitionsforschung für die Darstellung räumlichen Wissens? In: Freksa, C. & Habel C. (eds.) Repräsentation und Verarbeitung räumlichen Wissens. Springer-Verlag, Berlin.

Güsgen, H. W., 1989. Spatial reasoning based on Allen's temporal logic. ICSI TR-89-049, International Computer Science Institute, Berkeley.

Hernández, D., 1990. Relative representation of spatial knowledge: the 2-D case. To appear in: D.M. Mark & A.U. Frank (eds.) Cognitive and Linguistic Aspects of Geographic Space. Kluwer, Dordrecht.

Lakoff, G., 1987. Women, fire, and dangerous things: what categories reveal about the mind. University of Chicago Press, Chicago.

Lakoff, G. & Johnson, M., 1980. Metaphors we live by. University of Chicago Press, Chicago.

Reiter, R., Mackworth, A., 1989. A logical framework for depiction and image interpretation. Artificial Intelligence 41, 125-155.