In P. Burrough & A. Frank, *Geographic objects with indeterminate boundaries,* 109-121, London: Taylor & Francis 1996.

On the Relation between Spatial Concepts and Geographic Objects

Christian Freksa and Thomas Barkowsky
Fachbereich Informatik
Universität Hamburg
Vogt-Kölln-Straße 30
22527 Hamburg
Germany

{freksa, 1barkows}@informatik.uni-hamburg.de

Abstract

In processing knowledge about spatial situations, spatial concepts are employed for representing objects, properties, and relations in the world. This article presents some fundamental difficulties encountered in representing and processing knowledge about the real geographic world and motivates the need for sophisticated conceptual structures for dealing with spatial knowledge. Consequently, semantic and structural aspects of spatial concepts are discussed. Specific attention is paid to relations among different concepts on one hand and to the relation between conceptual structures and structures in the real world on the other hand. Issues like discrete vs. continuous, crisp vs. fuzzy, fine vs. coarse, and top-down vs. bottom-up concept formation are discussed in the context of spatial representations. The article derives suggestions for task-specific concept formation for the use in future Geographic Information Systems. It concludes with a discussion about fuzzy boundaries of geographic objects.

1 Why is it Difficult to Represent Geographic Knowledge?

The geographic world surrounding us is extremely complex. When we want to master a given problem in this world, we need to single out particular aspects of current interest from this multifaceted formation. So at any given time we are only interested in few objects, and concerning these objects again we are regarding only particular properties and/or relations. The capability of isolating the relevant aspects and relating them to one another, results in a unique intellectual efficiency. This efficiency, however, is necessary for successfully operating in the world.

To represent knowledge about the world is to make explicit specific aspects of the world. In making explicit certain aspects we ignore others. In representing knowledge, every single aspect of interest can be represented separately. Alternatively, we can aim at representing different aspects within a single structure; this requires that the different aspects of interest must be compatible, i.e., they must fit into one reference system corresponding to a global view. It is impossible to make all potentially interesting aspects of the world simultaneously explicit within one representation medium – be it a conventional map or a single representation structure in a computer.

1.1 The need for different world views

There is an information-structural reason why a unified representation of the geographic world is not sufficient for solving all the tasks that can be solved with that information: we are dealing with geographic objects and with relations between these objects. If we could decide once and for all which entities we should view as objects and which as relations, our task would be simpler; however, an important feature of using world knowledge intelligently is the ability to switch between views; depending on the specific task to be solved, certain entities may be viewed as objects (fixed background entities) and others become the relations we manipulate. For other tasks, these roles may change.

For example, we may consider roads as relevant geographic objects in some context. For solving certain navigation or transport tasks, we may be interested in the intersections between these roads; having roads as objects, intersections naturally can be viewed as (connection-) relations between roads. For designing road systems, it may be convenient to view road intersections as the primary objects; the roads then can be viewed as (connection-) relations between these intersections. Of course we could say, we want to view both, roads and their intersections, as objects to obtain a unified representation (after all, our toy train systems contain both, regular tracks and switches as objects). But this does not really solve the general problem: when we view both, the roads and their intersections as (separate) objects, we create a situation in which the roads do not meet the intersecting roads; all the roads meet intersections. To solve navigation tasks, for example, we will have to consider relations between roads and intersections; thus we only have shifted the problem to the next level. We encounter similar situations when we consider regions and their boundaries, which may be viewed as objects and relations, respectively, or vice versa.

Therefore, if we want to make all potentially interesting knowledge accessible, we must make it explicit in different structures (maps or computer representations). Creating many different structures for representing knowledge about the same domain becomes expensive – both computationally and in terms of storage, since a lot of implicit information must be carried along to link the knowledge to its domain.

1.2 The GIS forms the mediating instance between world and user

When developing a Geographic Information System (GIS), we must find an adequate compromise between two extreme possibilities: (1) acquiring and storing all knowledge from raw information once and for all before the knowledge is accessed, and (2) providing unprocessed raw information and computing specific knowledge on demand. On one hand we expect a GIS to contain sufficient data about the geographic world, on the other hand we want to obtain a selective view of the relevant aspects and to hide any other data of lesser interest. The entities represented in a GIS stand for the real world objects and their properties. From this point of view, GISs form the mediating instance between world's reality and the way humans interact with this reality. The user expects the entities represented in the system to show the same properties as the real objects they are standing for.

1.3 Focus on human use of spatial knowledge

The human capability of seeing the world as an inexhaustible origin of information may make it desirable to process the knowledge about the world in such a way that it is instantly available to the human user. But if we process this knowledge about the world in advance, we will create many structures which most likely will never be accessed. On the other hand, if we do not provide structured knowledge to the user, great efforts may be required to compute this knowledge when needed.

When we consider human capabilities of dealing with geographic information we can identify two challenges for the development of "intelligent" geographic information systems: (1) How can we model the human ability to focus on relevant information when solving a problem? and (2) How can we overcome the problem that a GIS can not contain all facts about the real world that might become important in a special context?

2 What are Spatial Concepts?

When we consider real world objects, we usually are interested in certain properties of these objects, i.e., we regard the objects under certain aspects. For example, when we take a look at a geographic entity, say a lake, we regard it with respect to horizontal extension, depth, shape, or the like. All these notions that describe spatial aspects of a subset of the world, we call spatial concepts. We will use the term 'concept' in a rather general sense (c.f. [Church 1956]). Concepts can be anything we have a notion of; thus, "size" can be a concept and "big" can be a concept as well. As we predicate aspects of objects using concepts, it is obvious that spatial concepts will play a crucial role for representing knowledge about the geographic world.

2.1 Properties of spatial concepts

In the following subsections we will present several semantic and structural aspects of spatial concepts which are of particular importance for their representation and for the operations that are to be performed on them. Specifically, we will address the relationships between different concepts and the relationships between concepts and real world entities.

2.1.1 Concepts have meaning wrt. objects

As we have already pointed out, we use concepts to describe aspects or properties of objects. This means, concepts are related to the aspects of the objects that are described. It is impossible to describe qualities of objects without using related concepts. Conversely, concepts have their meaning rooted in their relation to objects, e.g. the concept of a square is related to quadrilateral objects whose sides have equal length and meet at right angles. The mutual relatedness between concepts and objects enforces certain structures upon the concepts; in particular, not every relation between concepts and objects is meaningful.

2.1.2 Concepts have meaning wrt. other concepts

Nevertheless, the meaning of spatial concepts is not only given by their relations to physical objects. The meaning of concepts is given to a large extent by their relations to other concepts (c.f. "lateral thinking" [de Bono 1969]). We illustrate this point using the example of the square: we can view the meaning of the concept "square" by its relation to all existing physical objects of square shape; but we also can view the meaning of "square" by relating it to imagined or imaginable objects of square shape. Rather than viewing the meaning to be made up of an infinite number of relationships with imaginable objects, we can view the meaning to consist of a finite number of relationships between related concepts.

For example, the concept "square" may be related to the concept "rectangle" by a "special_case" relationship, to the concept "equilateral triangle" and "rhombus" by a "near_miss" relationship, to the concept "shape" by an "isa" relationship, etc. In this way, spatial concepts may have a meaning without reference to physical instances. In particular, spatial concepts have a meaning independent of an envisioned physical materialization: the shape concept "square" abstracts from the instantiation by a section of a checkers board or a market place in a city.

2.1.3 Concepts induce distinctions

Concepts are human made entities. We do not have any concept which is not related in some way to another concept. We form new concepts in order to make distinctions between features (or sets of features) in the world. Thus, concepts come in pairs or n-tuples, depending on how many cases are to be distinguished (c.f. [Freud 1910], [Quillian 1968]). The relation between concepts frequently can be viewed as positive, when there are concepts that describe similar aspects (e.g. "special–case"), or as negative when the description within one concept excludes others "square" vs. (general) "rectangle". Interestingly, we frequently find both, relations with positive and relations with negative connotation between a given pair of concepts (e.g., a square is a special case of a rectangle (positive) and it is not just any rectangle (negative). This fact about human concepts plays a particularly important role for constructing appropriate models for dealing with such concepts. The property of relatedness between concepts is often important for gaining implicit knowledge about conceptual aspects from the description of other concepts.

Consider a spatial concept like “big”. In attempting to define the meaning of “big” people sometimes ask questions like ‘how big is “big”?’ [Denofsky 1976]. Taking into account the foregoing considerations, we can answer this question on two levels: (1) When we view the meaning of concepts as given by instances, we can provide the answer “it depends”; specifically, it depends on the type of entity to which the concept “big” is applied. By applying the notion “big” to some specific instance or instances, we express a distinction from instances which are not “big”. If everything in the set of consideration was equal in size, it would not make sense to apply the concept. (2) When we view the meaning of concepts as given by their relations to other concepts, we will find a qualification relation to the concept “size” and a contrast relation to the concept “small” and possibly to the concept “very big”.

On level (2), the meaning of the concept “big” is independent of specific sizes to be distinguished, it is even independent of the types of entities it may apply to. If we hear about a “big” something and do not know which something or which type of something is being referred to, we still can answer questions about the size of the entity under discussion; specifically, we will be able to say that the size of the something is greater than the size of a corresponding something labeled “small” and that the size of a “big” something else of the same category will be of the same order of magnitude as the size of the something. Thus, there is a structural aspect of size concepts which is context independent; in essence, it establishes universal conceptual relations valid for all situations in which the concept can be used.

For optimal informativity it is desirable to make available a dynamically adaptable number of concepts, depending on how many objects are to be classified and how significantly they differ from each other [Freksa 1981]. For example, when there are only few sizes to be distinguished, say five, it is sufficient to differentiate five classes by making available appropriate concepts. In natural language, we use contextual adaptivity of conceptual classes all the time: depending on the situation, we implicitly contrast “big” vs. “small” or “big” vs. “very big” vs. “rather big” vs. “medium-sized” vs. … . Thus, the meaning of spatial concepts is determined by at least three types of relations: (1) relations between concepts and physical objects, (2) relations between concepts and related concepts (context of discourse), (3) relations between objects and situation context.

2.1.4 Concepts can be substituted for by finer or coarser concepts

It is one of the most important features in dealing with concepts that the number of classes they are generating can be increased or decreased just as needed for adequate description. This means, the spectrum of possible values is reorganized according to the situation’s needs. Of course, the single classes need not be equal in size, neither wrt. the range they cover nor wrt. the number of actual values they include. Figure 1 illustrates this property. It shows that a range associated with a concept, e.g. ‘size’ is further and further divisible, up to the refinement level actually needed, or conversely, fine concepts can be merged to form coarser concepts. This property requires a neighborhood structure for concepts: ‘fine’ concepts which are in opposition on a level of high resolution are united in harmony on a ‘coarser’ conceptual level on which fewer distinctions are made [Freksa 1991]. If we assume that new concepts result from the refinement of existing concepts, we will always obtain a neighborhood structure as depicted in Figure 1.

2.1.5 Concepts have no internal structure

We can view concepts as atomic discrete entities in the sense that we either use the concept or we don’t use the concept in a given situation. Concepts do not have an internal structure but concepts can be replaced by a structure of finer concepts or they can be merged to form coarser concepts. It is important to see that if we replace a concept by a structure of finer concepts, the original concept is no longer accessible; therefore we consider concepts as not internally structured.

When we replace a concept by a structure of finer concepts, we can observe interesting polarity properties of the competing finer concepts: on one hand, the finer concepts are *in opposition* to one another: we introduce the concepts “big” and “small” in order to emphasize opposing qualities; on the other hand, the new concepts are *neighbors;* they can be merged, when this particular distinction is not required. This dualistic relationship between concepts with common roots is extremely useful for generating adequate descriptions in different contexts and for interpreting them in a robust way.

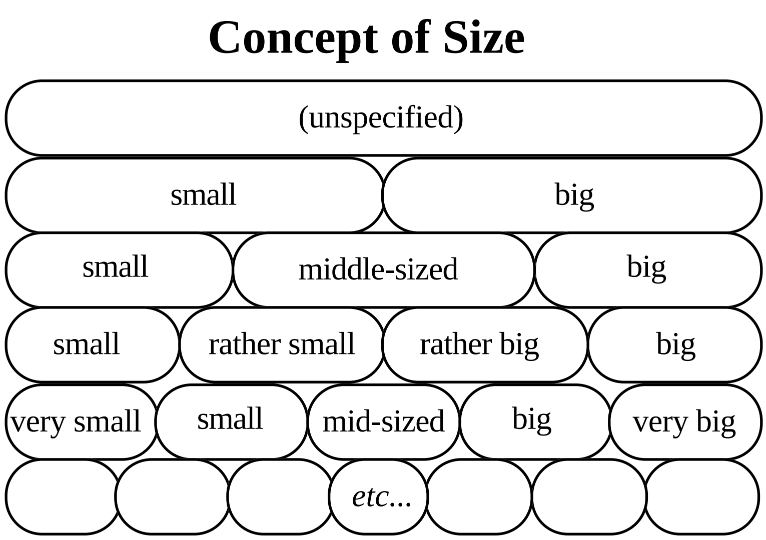

Figure 1: Substitution of Concepts

By stating that concepts have no internal structure we imply that it is not meaningful to speak about continuous or discrete concepts. If we assume that concepts are created by refining existing concepts, we do not gain anything by postulating a continuous (as opposed to discrete) internal structure. The assumption of continuous structures results from a world view which is not based on conceptual structures or on observed feature values but on a generalization of low-level features. This world view is heavily influenced by the availability of mathematical models which characterize complex structures in terms of simpler components.

2.1.6 Concepts are crisp wrt. their explicit competitors

An interesting question is if concepts are *crisp* or *fuzzy* entities, i.e., if they have sharp or gradual boundaries. To answer this question we must determine what the boundaries of concepts are. If we consider the relation between a given concept and an opposing concept, we clearly have a sharp boundary, as opposing concepts result from distinctions and distinctions are made on the basis of qualitative criteria rather than on the basis of (arbitrary) gradual values. For example, when we distinguish between “very big” and “big” we want to point out a *significant* difference in size.

We can now look at concepts which coexist on different branches of a ‘concept hierarchy’ as result of distinctions made in different situations rather than on the basis of an opposition as described above. For example, we may have notions like “pretty big” and “rather big” which do not seem to be in opposition but which are not identical either. Do the associated concepts have crisp or fuzzy boundaries to one another? Well, if these concepts are on non-neighboring branches of our conceptual hierarchy, they may share no boundaries; thus, the question may be ill-posed. “However”, you might argue, “these concepts address the same feature dimension, they express almost identical ideas, so what is their relationship to one another?”

If we continue to restrict our considerations to the conceptual domain, the relationship between concepts on different branches of the conceptual hierarchy is given by the vertical relationships which connect these concepts; but these relationships may not induce boundaries between the concepts in question. However, if we take into consideration the relationships between spatial concepts and objects in the real or imagined world, the situation changes: when asking the question of applicability of concepts in a concrete situation, the question of fuzzy vs. crisp boundaries of concepts appears in a new light.

2.2 Relations between concepts and objects

In section 2.1 we have discussed spatial concepts and their properties in isolation from the spatial domain they correspond to. In the present section we will address aspects pertaining to the correspondence between conceptual structures and real-world entities (c.f. [Palmer 1978]).

2.2.1 Lack of detail results in fuzziness

Fuzziness is not a property of a single aspect or concept, but rather the property of a relation between a concept and an object. When concepts come into existence 'top-down' by refinement of existing concepts rather than 'bottom-up' by definitions in terms of low-level features, then the relation between concepts and specific object properties may be viewed as 'fuzzy' as we may not be able to determine for each low-level value if the relation holds or if it does not. In a sense, we apply coarse notions to situations of fine resolution; the neighborhood structure among concepts and domain situations allows for interpolating and extrapolating between central uses of concepts and clearly inapplicable situations. The result is a fuzzy correspondence. Admitting fuzzy correspondence in representations is very useful, as it enables us to place situations into ballparks of closely matching conceptualizations when no precisely matching conceptualization is possible or desirable.

2.2.2 Horizontal and vertical neighborhood

We can distinguish two dimensions of neighborhood between concepts: 'horizontal' and 'vertical' neighborhood. 'Horizontal neighborhood' refers to competing concepts on the same level of granularity [Freksa & López de Mántaras 1982] while 'vertical neighborhood' refers to compatible concepts on different levels of granularity (compare Fig. 1) [Zadeh 1978, Hobbs 1985]. Thus, exploitation of the horizontal dimension allows for the selection of an appropriate description *value* while exploitation of the vertical dimension allows for the selection of an appropriate description *granularity*, depending on the context of the specific situation.

The combination of the horizontal and vertical degrees of freedom allows for great flexibility in selecting the most appropriate descriptor, in a given situation. For example, suppose coarse granularity is sufficient in a given situation (e.g. 'small', 'middle-sized', 'big' – compare Fig. 1), but the object to be characterized is equally well described by two competing descriptors (e.g. 'small' and 'middle-sized'), i.e., the situation is a boundary case; then we can go to the next finer level of granularity and typically we will find a concept which very well characterizes the object (e.g. 'rather small'). In this way, 'fuzziness' of concepts can become instrumental in selecting appropriate descriptors – which then may be in a rather unfuzzy (crisp) relation to the described situation.

2.2.3 Concept hierarchies do not form trees

In refining concept spaces, we do not simply subdivide individual concepts as in a hierarchical tree (c.f. [Winston 1975]), but we revise the subdivisions according to the neighborhood structure in such a way that the boundary cases on one level of description become central cases on the next level. In other words, when a concept is needed to refer to a boundary region of two (relatively coarse) concepts, the new (finer) concept typically refers to a region integrating both sides of that boundary (compare Fig. 1). By exploiting the horizontal and the vertical dimension of concept hierarchies, we obtain a powerful tool for tailoring effective and efficient descriptions of arbitrary situations from arbitrary points of view.

3 How do we Describe the World with Concepts?

In this section we address the problem of structural incompatibilities between the conceptual domain and the object domain. The structure of the object domain depends largely on the view taken with respect to this domain: when we view the spatial domain as *dense* or *continuous*, we implicitly refer to the set of *possible situations* rather than to actual situations. If we take each given situation individually by itself, we do not obtain a continuous structure. Thus, the structural compatibility between concepts and objects depends much on the assumed world view.

3.1 The properties of the world are assumed to be continuous

When we regard an entity in the world under a certain aspect, say a lake under the aspect ‘area’, we assume that *in principle* lakes may have any extension, within a certain range. This insight is responsible for using a continuous scale, i.e. real numerical values, to describe areas. In fact, this is a rather sensible approach to data acquisition.

When concerned with special tasks to be performed on this data, the pure numerical description may not be of major interest. What we actually need is a classification of the data with respect to the task of current interest. Each individual task typically induces specific thresholds, i.e. critical values above or below which there may be a qualitatively different situation. These thresholds may be crisp or fuzzy.

3.2 Concepts are discrete entities

As suggested in section 2, new concepts are formed when new distinctions are needed; thus concepts – like the symbols denoting them – are inherently discrete entities. How do we use these discrete entities to describe a world whose objects are assumed to have continuous-valued properties? To answer this question we will look at the formation of concepts from classifying real-world entities on one hand and from refining existing concepts on the other hand.

We have seen that the world of real objects and their properties may appear quite different from the world of concepts we use for describing the ‘real world’. When we describe certain aspects about the ‘real world’, we have to match the (discrete) world of concepts with the (continuously perceived) world of features of the entities in the real world. To do this, we can (1) go ‘bottom-up’ from the features towards the concepts, (2) go ‘top-down’ from the concepts towards the features, or (3) combine the two directions. At some point, we will reach a level where the bottom-up generalization of feature values will reach about the same granularity as the top-down refinement of the concepts. Nevertheless, we cannot expect a one-to-one mapping between the top-down conceptualization and the bottom-up generalization; rather, there will be incompatibilities. This phenomenon is illustrated in Fig. 2.

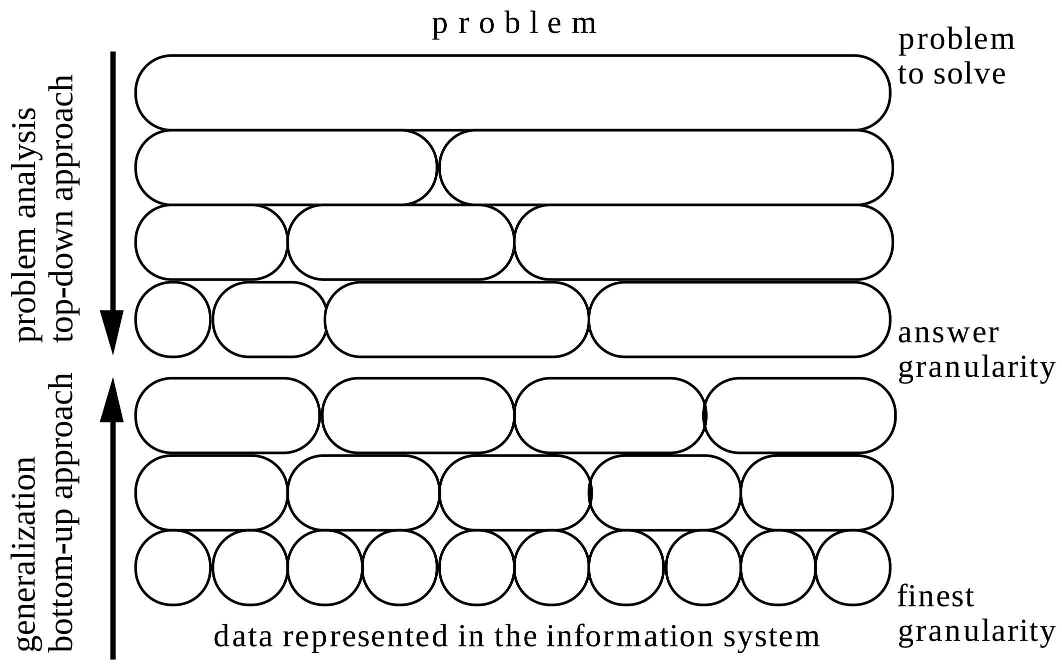

Fig. 2: Bottom-up vs. top-down conceptualization

3.3 Concepts are adjusted to the situation, not the other way around

As the discussion of the incompatibility between the top-down and the bottom-up conceptualization shows, it does not make sense to form concepts mechanically from the real world data alone. This is, because there is no absolute measurement for an entity to belong to one class of a concept or to another. The same set of data may require different classifications (even at the same level of granularity) when used within different contexts.

Concepts provide an efficient means for understanding, thinking, planning and communicating. But concepts are not fixed in their reference to real world entities; rather they allow to be adapted with some flexibility to the requirements of the specific situation. By using concepts we can see the world as detailed as we need, or, by using concepts we construct our reality of interest.

3.4 GISs eventually should employ concept-driven sensors

The epistemological independence between concepts and real world entities can be used when we are concerned with the level of granularity to be chosen for the representation of data in technical systems. Let us assume, we do not have acquired all data we need to answer a certain question in a system (this is a rather realistic assumption). Let us further assume, that we have implemented an appropriate model of human conceptualization. Then it is quite imaginable to combine a technical system with an automatic information acquisition system (e.g. satellite-based) to obtain the data needed according to the conceptualization wanted directly from the real geographic world.

The task-driven application of conceptualizations to knowledge acquisition would be a very appropriate compromise between the two options described at the outset of the paper: (1) acquiring and storing all knowledge from raw information once and for all before the knowledge is accessed, and (2) providing unprocessed raw information and computing specific knowledge on demand. The compromise consists of preprocessing raw data to a level of conceptually relevant granularity and using this preprocessed data on demand by the conceptualizations formed in specific task contexts.

The suggested approach exhibits a strong correspondence to the human use of geographic maps. A geographic map corresponds to preprocessed raw data: certain features are classified according to criteria from the conceptual level, other features – in particular spatial relations – are only processed with respect to data resolution (granularity). Depending on the task to be performed, the conceptual level determines, which relations in the map are further elicited or processed.

4 Where is the Fuzziness in Geographic Objects?

In representing knowledge about the geographic world, we must distinguish different kinds of uncertainty: 1) We may not know the precise location of crisply classified geographic entities and thus may be uncertain about their location, or 2) we know the precise locations of the geographic entities including the – possibly gradual – transitions between them, but we are uncertain precisely how to classify them, or 3) we may have a combination of the two possibilities.

4.1 Uncertainty due to imprecise knowledge

In the first case, we are confronted with the problem of representing and interpreting incomplete information. Information may be incomplete due to a variety of reasons: for example, a piece of information may be associated with a European metropolis and we may be uncertain if London or Paris is meant; thus, we have two *discrete* possibilities with regard to the correct reference and the information is incomplete with respect to the resolution of this ambiguity. On the other hand, we may be uncertain with regard to the precise location of a geographic object due to measurement errors; here, we obtain a *continuum* of possibilities with regard to the precise location of the object and the information may be considered incomplete with respect to the required precision.

Discrete possibilities frequently are represented by logical disjunctions and interpreted like an 'exclusive or': either London or Paris (but not both). We obtain a special case of the discrete situation when the different possibilities are arranged according to some ordering criterion. Under certain circumstances (which are given in many situations), the possibilities are distributed in such a way that only *neighborhoods* of possibilities can occur. These possibilities can also be represented by disjunctions: for example, in specifying geographic locations, “38 meters” may mean in more precise terms “37.5 meters” or “37.6 meters” or … “38.4 meters”.

Treating neighboring possibilities in this way has computational drawbacks (especially computational complexity) and it neglects the spatiality expressed in the neighborhood of the geographic locations as it ignores the difference between close and remote locations. Therefore, representing neighboring possibilities by structures preserving neighborhood has great advantages. Continuous structures preserve neighborhood information. Intervals, probability distributions, and possibility distributions are such structures. However, in dealing with real situations, most often we have discrete information to begin with; therefore there are definite advantages to employ discrete models which preserve the ordering information relevant in the domain.

4.2 Uncertainty due to indefinite correspondence between concepts and represented world

In the second case (precise information, unclear classification) we are confronted with a rather different problem: the uncertainty is not associated with the information itself but it is associated with the correspondence between the concepts available for the description of the geographic world and the geographic world itself (c.f. [Rosch 1973, 1975]). If we use as many concepts as possibilities we want to be able to distinguish and use a one-to-one correspondence, the problem has vanished. However, if these concepts are not connected appropriately, we loose again the internal structure of the domain and the representation is not very informative. What we usually want is to represent knowledge in such a way that few concepts capture rich situations. In particular, this means that (1) gradual transitions from one concept to another should not be captured by introducing a large number of intermediate concepts and (2) we need concepts that are adaptable to a multitude of situations if the concepts are to be used universally.

We will address here specifically the representation of transitions. The classical approach to representing gradual transitions is the use of fuzzy sets [Zadeh 1965]. A main utility of the fuzzy set approach is that concepts are split into two levels of consideration: (1) the level on which we talk about concepts (the symbol level) and (2) the level on which the correspondence to real world entities is established (the semantic level). In this way, we obtain an interface between the discrete world of concepts and a possibly continuous world of entities to be described [Freksa, in print]. The fuzzy set approach assumes that the concepts associated with the labels of fuzzy sets can be defined precisely. In everyday situations, in particular in non-technical domains, this assumption is usually too strong.

A weaker – and in many situations more realistic – assumption is that we can very well structure the concepts we use and that we can very well structure the entities and features we want to capture, but we have difficulties in establishing a precise correspondence between concepts and real world entities. The reference of the concepts then can be established by structure matching rather than by correspondence between individual conceptual and real world entities. As a result of this process, a certain fuzziness in the correspondence between concepts and real world entities persists: we may have several possibilities in assigning concepts to a given situation. This fuzziness is benign, as the applicable concepts will be very closely related alternatives due to the preservation of the structure.

5 Conclusion

We have addressed the problem of representing knowledge about the real world in general and representing spatial knowledge about the geographic world in particular. The focus of the discussion was placed on issues related to boundaries in the represented world, in the representing world, and in the relation between these two worlds. We face the challenge of dealing with (1) complex open worlds, i.e. worlds whose dimensions and values cannot be entirely specified; (2) incomplete and imprecise information, i.e. information which cannot be taken at 'face value'; (3) knowledge and concepts of varying granularity causing fuzzy correspondence between concepts and real world entities; (4) knowledge which is to be used for different tasks requiring different resolution and different conceptualization; (5) knowledge in different contexts requiring different discrimination capabilities.

We have investigated the role of concepts and conceptual structures in representing knowledge about the real world and emphasized their autonomy with respect to the entities and assumed structures in the represented domain. We have argued that geographic objects are not fuzzy in themselves, but that they are fuzzy with respect to the precision of the underlying knowledge and/or with respect to their classifiability in terms of a given set of concepts. In order to deal with gradual boundaries in a sophisticated way, we propose an approach which takes into account 1) the neighborhood structure of geographic entities according to a physical model, 2) the vertical and horizontal neighborhood structure of the spatial concepts, 3) the correspondence between concepts and geographic entities.

The presented work is based on studies of representing soft knowledge for the purpose of human-human and human-machine communication about open worlds and on formal approaches to representing qualitative temporal and spatial knowledge.

References

de Bono, E., 1969, *The mechanisms of mind*, New York: Simon & Schuster.

Church, A., 1956, *Introduction to mathematical logic, I,* Princeton, N.J.: Princeton University Press.

Denofsky, M.E., 1976, How near is near?, *AI Memo No. 344*, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge.

Freksa, C., 1981, Linguistic pattern characterization and analysis, Ph.D. thesis, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley. Abstract in: *J. of Pragmatics* 6 (1982), 371-372.

Freksa, C., López de Mántaras, R., 1982, An adaptive computer system for linguistic categorization of "soft" observations in expert systems and in the social sciences, in: *Proc. 2nd World Conf. on Mathematics at the service of man*, Las Palmas.

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

Freksa, C., in print, Fuzzy Systems in AI, in Kruse, R. et al. (eds.) *Fuzzy Systems in Computer Science*, Wiesbaden: Vieweg.

Freud, S., 1910, Über den Gegensinn der Urworte, *Jb. psychoanalyt. psychopath. Forsch.,* Bd. 2 (1), 179-184.

Hobbs, J.R., 1985, Granularity, *Proc. 9th Intern. Joint Conf. on Artificial Intelligence*, pp. 432-435.

Palmer, S.E., 1978, Fundamental aspects of cognitive representation, in Rosch, E., Lloyd, B. (eds.), *Cognition and categorization,* Hillsdale: Erlbaum.

Quillian, M.R., 1968, Semantic memory, in Minsky, M. (ed.), *Semantic information processing,* Cambridge: MIT Press.

Rosch, E., 1973, On the internal structure of perceptual and semantic categories, in Moore, T.E. (ed.), *Cognitive development and the acquisition of language*, New York: Academic Press.

Rosch, E., 1975, Cognitive representations of semantic categories, *Journal of Experimental Psychology: General, 104.* 192-233.

Winston, P.H., 1975, Learning structural descriptions from examples, in Winston, P.H. (ed.)*The psychology of computer vision,* New York: McGraw Hill.

Zadeh, L.A., 1965, Fuzzy sets, *Information and Control* 8, 338-353, .

Zadeh, L.A., 1977, Fuzzy sets and their application to pattern classification and clustering analysis, in van Ryzin, J. (ed.), *Classification and clustering,* New York: Academic Press.

Zadeh, L.A., 1979, Fuzzy sets and information granularity, in Gupta, M.M., Ragade, R.K., Yager, R.R. (eds.), *Advances in fuzzy set theory and applications,* Amsterdam: North-Holland.