

Qualitative Spatial Reasoning Using Orientation, Distance, and Path Knowledge*

Kai Zimmermann and Christian Freksa

Department of Computer Science, University of Hamburg
Bodenstedtstr. 16, 22765 Hamburg, Germany
zimmerma@informatik.uni-hamburg.de
freksa@informatik.uni-hamburg.de

Abstract We give an overview of an approach to qualitative spatial reasoning which is based on directional orientation information as available through perception processes or natural language descriptions. Qualitative orientations in 2-dimensional space are given by the relation between a point and a vector. The paper presents our basic iconic notation for spatial orientation relations which exploits the spatial structure of the domain and explores a variety of ways in which these relations can be manipulated and combined for spatial reasoning. Using this notation, we explore a method for exploiting interactions between space and movement in this space for enhancing the inferential power. Finally the orientation based approach is extended by distance information, which can be mapped into position constraints and vice versa.

1 Introduction

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 through the other channels. As a consequence, we are disproportionately confident about what we know about space: we take it for real

Our research on spatial representations and reasoning is motivated by the intuition that ‘dealing with space’ should be viewed as cognitively more fundamental than abstract reasoning. Afterall, one of the very first tasks we learn to accomplish is to orient ourselves in the environment. The use of spatial metaphors in language and problem solving tasks also indicates that there might be a specialised, maybe less expressive, but optimized, spatial inference mechanism. Why else would we translate a problem into the specialised domain of space if the domain of space is handled by a general inference mechanism? As a consequence, we want to understand dedicated spatial reasoning before we construct general abstract reasoning engines. The goal of this research is the conception of a ‘spatial inference engine’ which deals with spatial knowledge in a way more similar to biological systems than systems based on abstract logic languages.

Spatial information, or more specifically, directional information about the environment, is directly available to animals and human beings through perception, and is crucial for establishing spatial

* Support from the DAAD PROCOPE program and from the HANSISCHE UNIVERSITÄTSSTIFTUNG is gratefully acknowledged

location and for path finding. Distance information is directly available, too, when take into account the concept of motion. Such information typically is imprecise, partial, and subjective, but the more we explore the environment the better our knowledge about it gets, i.e. there must exist a mechanism to combine and to integrate multiple observations into a representation with increasing granularity. In order to deal with this kind of spatial information we need methods for adequately representing and processing the knowledge involved. In this paper we present an approach for representing and processing qualitative spatial information which is motivated by cognitive considerations about the knowledge acquisition process. The approach includes ways for dealing with orientation, position, motion, and distance information.

Consider a simple localization task: you walk straight along a road, turn to the right, walk straight, turn left, and walk straight again. Now you would like to know where you are located with respect to the first road you walked on. Tasks like this are very fundamental for almost all animals and human beings. We mostly carry them out subconsciously – except when we fear to get lost, for example in underground walkways. In the following we describe how we represent this knowledge for modelling spatial reasoning.

2 Overview of Existing Approaches

A variety of approaches to qualitative spatial reasoning has been proposed.. Gsgen [1989] adapted Allen's [1983] qualitative temporal reasoning approach to the spatial domain by aggregating multiple dimensions into a Cartesian framework. Gsgen's approach is straightforward but it fails to adequately capture the spatial interrelationships between the individual coordinates. The approach has a severe limitation: only rectangular objects aligned with their Cartesian reference frame can be represented in this scheme. Since we only represent the relative position and orientation information of points we are not restricted to one specific rectangular coordinate system that has to be applied to all objects.

Cui, Cohn and Randell [1993] attack the problem of representing qualitative relationships involving concave objects. They introduce a 'cling film' function for generating convex hulls of concave objects; they then list all qualitatively different relations between an object containing at most one concavity and a convex object. Egenhofer and Franzosa [1991] develop a formal approach to describe spatial relations between point sets in terms of the intersections of their boundaries and interiors. They do not use orientation information.

Hernndez [1992] considers 2-dimensional projections of 3-dimensional spatial scenes. He overcomes some deficiencies of Gsgen's approach by introducing 'projection' and 'orientation' relations. For the dimension of projection he adopts and extends the ideas of Egenhofer [1989], i.e. the binary topological relationship between two areas in the plane. But he combines the topological information with relative orientation information that can be defined on multiple levels of granularity. Nevertheless, he is still describing scenes within a static reference system. Freksa [1991] suggests a perception-based approach to qualitative spatial reasoning; a major goal of this approach is to find a natural and efficient way for dealing with incomplete and fuzzy knowledge.

Schlieder [1990] develops an approach which is not based on the relation between extended objects or connected point sets. Schlieder investigates the properties of projections from 2-D to 1-D and specifies the requirements for qualitatively reconstructing the 2-dimensional scene from a set of projections yielding partial arrangement information.

Frank [1991] discusses the use of orientation grids ('cardinal directions') for spatial reasoning. The investigated approaches yield approximate results, but the degree of precision is not easily controlled. Mukerjee and Joe [1990] present a truly qualitative approach to higher-dimensional

spatial reasoning about oriented objects. Orientation and rectangular extension of the objects are used to define their reference frames.

3 The Representation

3.1 Motivation

Although a lot of formalisms for spatial reasoning do already exist they do not deal with large scale navigation or they do not appeal from a cognitive point of view. Our approach is motivated by cognitive considerations about the availability of spatial information through perception processes, see Freksa [1991]. A major goal of this approach is to find a natural and efficient way for dealing with incomplete and fuzzy knowledge. Thus, a new representation has been developed with the following goals in mind:

- The representation should be simple and extendable.
- The formalism should allow for different levels of granularity, as well in the representation, e.g. if only imprecise knowledge is available, as in the choice of operations, e.g. under time constraints faster computation of partial results should be possible.
- The approach should resemble some fundamental properties known about human spatial reasoning to be plausible from a cognitive point of view.

One of the major differences to previous approaches is that the relative positions of other objects are not described with respect to (wrt.) one position but wrt. a vector that describes the movement between two positions. The operations applicable on this kind of representation are described below. Our representation allows us to describe orientation and position qualitatively, but it does not deal with the shapes of objects. Furthermore, in our formalism the operations do not yield approximate values but correct ranges of values. Other approaches by the authors with different base domains and entities have not yielded satisfactory results until we developed the representation scheme described in the following sections.

3.2 The Representation

Consider a person walking from some point a to point b . On his way he is observing point c . He wants to relate point c to the vector ab . For this he can, for example, make the qualitative distinction whether c is to the left or to the right of the line going through a and b . Given this line he can in addition ask whether c is before or behind a and b , respectively, when travelling along the vector ab . This kind of knowledge is easy to obtain while following a path or being at its end points. Thus he obtains a reference system that allows him to describe the position of c with increasing sharpness. We describe the situation in which he can distinguish 15 possible relations. If for some reasons it is not possible to decide whether c is behind or in front of b , for example, we end up with a disjunction of possible relations. See Fig. 1 for an example.

The obtained 15 relations form a conceptual neighborhood as defined in Freksa [1992a]. Note that it is not necessary to have the observer at point b . You can as well choose point a to be the standpoint of the observer who sees point b and c and relates the position of c to the line of sight to point b . In this kind of application of the formalism it might be harder to obtain the knowledge whether b or c is farther away, though.




Fig. 1 1) Consider somebody walking from a to b . On his way he observes c in front and to the left of b and d in front and to the right of b . 2) By introducing the two lines orthogonal wrt ab through a and b and the line through a and b we get an orientation grid with 15 different positions: six areas, seven places on the lines, and two points. 3) The positions of c and d can now be described in terms of these 15 spatial relations which is depicted iconically.

Although the chosen reference system defines a local orthogonal grid, the kind of information needed to conclude the relation between a point and a vector is easy to obtain. You can draw the distinctions between left and right, before or behind, at any time of the travel, each time increasing your knowledge. As it is known from experimental psychology that humans are poor at estimating angles and tend to use rectangular reference systems we think that the right angles we have based the formalism on are a good choice. We are not sure whether a finer degree of angular resolution is suitable from a cognitive point of view as the base of the representation, although this is possible in principle, see Ligozat [1993] for example. There are, however, means of describing the position of c with a higher degree of resolution in our formalism, if, e.g., the domain of distance is taken into account. Refer to section 8 for a detailed discussion.

4 Composition

Up to now we have presented a representation frame that allows us to specify the position of a point relative to a vector. We will now introduce two methods for composing these reference frames and to perform a constraint propagation in a network of relations. We call these methods *COARSE* and *FINE COMPOSITION*, respectively.

COMPOSITION is an operation defined on two relations $ab:c$ and $bc:d$ that yields the relation $ab:d$ as result. This operation allows us, for example, to traverse a path from a to b to c to d and to answer the question where we ended, i.e. point d , wrt. the first part of the path, i.e. vector ab , given only the partial knowledge $ab:c$ and $bc:d$. See Fig. 2 for an example.




Fig. 2 The for the path $abcd$ the COMPOSITION of the relations $ab:c$ and $bc:d$. The result is a disjunction meaning that d can be everywhere on the right of vector ab , but not on the line through a and b or to its left. The result can not be sharpened without further knowledge available, e.g. about the length of ab and bc or different paths.

4.1 Coarse Composition

COARSE COMPOSITION is an efficient generalization of the COMPOSITION operation. It combines neighborhoods of fine relations to form a coarse relation; the COMPOSITION then is carried out on the coarse relation. Typically, but not necessarily COARSE COMPOSITION leads to a coarser result. COARSE COMPOSITION only takes orientation knowledge into account, i.e. it deals only with the relative orientation of the vectors, but not with their length. See Fig. 3a for an example and refer to Freksa [1992b] for a detailed discussion.

4.2 Fine Composition

FINE COMPOSITION takes into account that due to the orthogonal lines through a and b there is a kind of rough distance knowledge available which can be exploited. Thus, for some combinations we can obtain better results. See Fig. 3b for an example and Freksa [1992b] for a detailed discussion.




Fig. 3 For the path $abed$ (Fig. 2) COARSE COMPOSITION of $ab:e$ and $be:d$ yields that d is somewhere behind the orthogonal line through b . The grayed dots describe the used coarse relation. With the FINE COMPOSITION we get the result that d is even behind the orthogonal line through a , because we know that e is behind b in the first relation.

Thus, we have two operations with different granularities from which we can choose according to the available resources. It should be noted, however, that although the operation of COARSE COMPOSITION can be executed faster than the FINE COMPOSITION operation it typically leads to longer constraint propagation time. This is because the chance of sharpening a relation obtained by COARSE COMPOSITION when combining it with results obtained via a different propagation path is higher than with the results of the FINE COMPOSITION operation, which leads to an additional

propagation of the sharpened results and a longer overall computation. The main advantage of COARSE COMPOSITION appears in situations where no fine relations are available or where a fine relation is subsumed by a coarse relation. Here, COARSE COMPOSITION can avoid the necessity of exploring disjunctive alternatives and thus prevent the problem of combinatorial explosion. See Fig. 4 for an example of how two different propagation paths can be combined. Although each step of the COMPOSITIONs leaves us with a disjunctive result, we end up with one single relation after combining the results



Fig. 4 Combining the results from both path *abcd* and *abed* of the above example, i.e. taking the intersection of the resulting relations, since both cases must be true, restricts *d* to be on the right, behind the orthogonal line through *b*, for COARSE COMPOSITION, or *a*, in the case of FINE COMPOSITION, respectively.

5 Additional Operations

Up to now we have presented the COMPOSITION operation which allows us to draw conclusions in the case of chaining paths. Now we will focus on operations that allow us to change the reference vector within one relation. With these additional operations we are able to compute the relation for every possible permutation of points. For a detailed discussion see Freksa and Zimmermann [1992]

5.1 Inversion

The first operation is called INVERSION (INV). It maps the relation between *ab:c* to the relation between *ba:c*, i.e. it inverts the orientation of the reference vector. See Fig. 5 for the exact mapping of the operation INV.




Fig. 5 The table shows the results of the INVERSION operation in the corresponding positions.

5.2 Homing

The next unary operation we will focus upon is called HOMING (HM). This operation maps the relation $ab:c$ to $bc:a$, i.e. we ask about where we have come from when further proceeding from b to c . See Fig. 6 for the results of this operation.




Fig. 6 The results of the HOMING operation in the corresponding positions.

Note that the HOMING operation allows us to subsume the qualitative navigation approach presented by Levitt et al. [1987], see Fig. 7. When standing at some given point and taking a panorama view we can determine our position relative to the axis through all points from the order in which the points occur in the panorama. If, for example, B appears on the right of A we can conclude that we are on the right side of the line running from A through B. This forms the basic source of knowledge in the approach proposed by Levitt et al. and can be modeled through the use of HOMING.




Fig. 7 The use of the operation HM to model the qualitative navigation approach proposed by Levitt et al.. Of course, if sharper knowledge about the position of B wrt. OA is available one gets better results for the position of O wrt. AB .

5.3 Shortcut

The last operation is called SHORTCUT. Given the relation $ab:c$ it yields $ac:b$, i.e. the position of b if we take the shortcut from a to b . See Fig. 8 for the results.




Fig. 8 The results of the SHORTCUT operation in their corresponding positions.

5.4 Example

With these operations we are able to compute relations for all possible combinations of points. We can now combine them to form other kinds of dual operations than COMPOSITION, e.g. we can compute the relative position of objects wrt. the next part of the path if their position to the current path is known. With this knowledge we can provide an agent with reassuring conditions that must be true when he is still on its correct way.

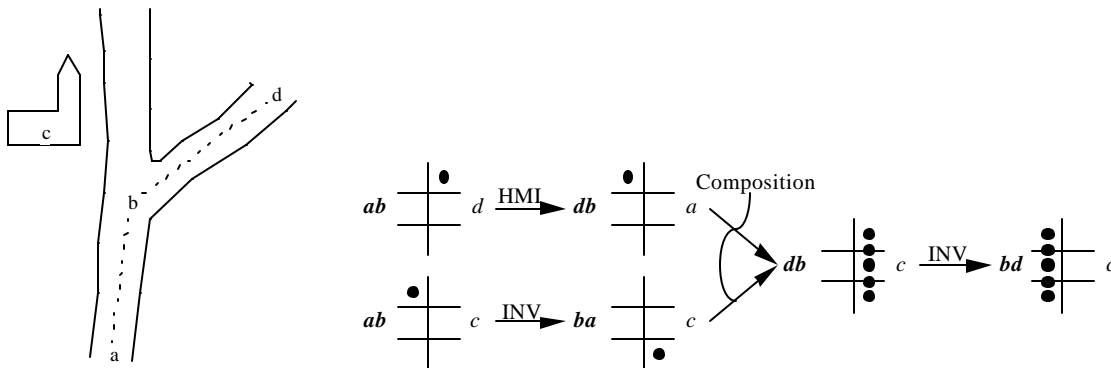


Fig. 9 The prediction of path assuring conditions. HMI is short for $INV(HM(x))$. When the positions of the points c and d is known wrt. ab we can predict the positions at which c must be wrt. bd , when proceeding further. Thus an agent can check whether it is still on its right path, i.e. bd , even if c becomes obstructed, e.g. by some obstacles.

6 Algebraic Combination of Operations

Fig. 10 shows how the operations can be combined algebraically. This kind of combination is not commutative, but it is associative. The associativity allows us, for example, to apply a general and possibly parallel constraint propagation algorithm in which the temporal order of combination does

not matter. If the combination were not associative, we would be restricted to an ordered computation, e.g., backward chaining.

O	ID	INV	SC	SCI	HM	HMI
ID	ID	INV	SC	SCI	HM	HMI
INV	INV	ID	HM	HMI	SC	SCI
SC	SC	SCI	ID	INV	HMI	HM
SCI	SCI	SC	HMI	HM	ID	INV
HM	HM	HMI	INV	ID	SCI	SC
HMI	HMI	HM	SCI	SC	INV	ID

Fig. 10 The algebraic combination of operations. The result of $HM(SC(x)) = HMI(x)$ can be found in the fourth row, column six. SCI is short for $INV(SC(x))$.

From this table we can see that HM and INV alone would be sufficient to generate all other operations, because $?(HM(x)) = SCI(x) = INV(SC(x))$, by applying $?? \vee \wedge$ we can produce SC, and so on. We provided the other operations since they have a natural meaning and they allow us to define complex manipulations more easily.

One problem is that the operations HM and SC sometimes yield a disjunction as result. It has been pointed out to the authors that this can be fixed when the underlying representation frame is extended by a circle with diameter ab , Fig. 11, see Latecki and Röhrig [1993] for details. Although this is a technical enhancement which resolves the two disjunctions there is no evidence that humans are capable of estimating whether an object is inside or outside that circle and it still does not fix the overall disjunction in point a and b .




Fig. 11 HOMING and SHORTCUT on a representation that has been extended by a circle with diameter ab . This fixes the disjunctions obtained at positions between the orthogonals through a and b but the universal relation in a and b still remains.

7 Using Path Knowledge

The representation of spatial orientation knowledge introduced above was originally designed for representing relationships between static positions of landmarks. We now introduce a dynamic component: motion. While in the representation described above, a single location was related to a reference vector, we now relate a motion sequence leading from the end point of the vector to that location. In the case that the relation consists of several possible locations we derive several possible paths. Thus, instead of reasoning about static situations, we take into account the possible motion sequences through the relation space according to the conceptual neighborhood structure, see Zimmermann and Freksa [1993] for details.

The representation consists of two levels: a disjunction of equally possible sequences and the underlying sequences themselves. Sequences are enclosed by square brackets and show the different intermediate states the mover will enter on his path. Although, the resulting sequences may seem trivial to a human observer, they capture knowledge about the structure of space that was not available before, since possible locations were just elements in an unordered set. The sequences are grouped via curly brackets and form an exclusive disjunction, i.e. only one of them may be chosen.

Example

In the static representation, the knowledge that c is on the right back wrt. vector ab is depicted by one relation, see Fig. 12. This representation is now transformed into the sequence of intermediate relations depicting the path from b to c . The underlying assumption is the direct connection of b and c by a straight line. This results in the sequence depicted below.




Fig. 12 The static representation is transformed into a sequence of intermediate states.

Imagine now that the person is walking down the street from a to b , then turns right at b and suddenly notices at c that there is a house on the right that had been occluded by trees previously (Fig. 13). Whereas in the former approach he could draw the conclusion about the static position of the house wrt. vector ab , the person is now able to derive knowledge about possible shortcuts from b to the house.


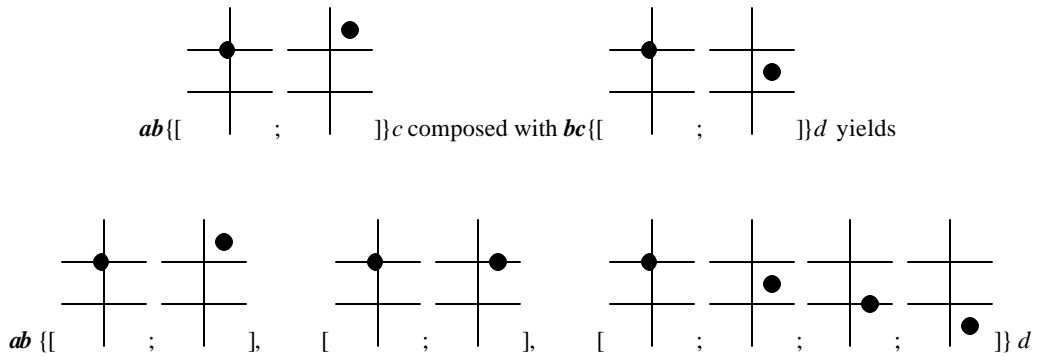


Fig. 13 A house occluded by trees on the first part of the path and the static result.

In the static approach, each black dot denotes a possible position of d related to ab . In the motion-based approach, we interpret the input of the calculation as descriptions of motions. Thus we obtain three possible sequences as result:



This means that if the person walks from point b to point d there are three possible qualitative directions a shortcut from b to d could have (see also Fig. 14):

- i) walk ahead to the right,
- ii) walk perpendicular to the right,
- iii) walk to the right back.

In the third case we are able to make predictions on his future encounters on his path, which may be used to guide his orientation about where to expect the house. To reach the last possible location of the house, e.g., he has to cross over the position of point a again, which would suggest to find a shortcut, not from b , but from a .




Fig. 14 The resulting possible sequences resolved by both, direction and time.

8 Adding Distance*

Up to now we have dealt with position and orientation in both the static and the dynamic approach. We will now show how knowledge about distances can be added to the representation. For a detailed discussion of this see Zimmermann [1993a] and for an introduction to λ -calculus, the underlying formalism used for enhanced distance reasoning, see Zimmermann [1993b].

In the above described reference frame three vectors occur explicitly: The vectors AB , BC and CD . These are now mapped from vectors to unoriented edges, since we want to exploit their distances. Additionally we introduce the orthogonal distance between point C and line AB , D_x , and the distances D_yA and D_yB between point C and the two orthogonal lines. See Fig. 15 for the resulting edges.




Fig. 15 The introduced edges. Edge a , b , and c coincide with the vectors BC , AC , and AB , correspondingly. The edges D_x , D_yA and D_yB decompose edges a and b orthogonally.

From these edges we take a further abstraction: their length. The lengths are represented symbolically and related via λ -calculus. Each kind of knowledge, length and position / orientation is treated separately by agenda based domain experts which communicate via a black board structure.

8.1 The Mapping Between Position and Distance Information

This section deals with how the different knowledge sources interfere. As we can see in Fig. 16, the distances restrict the possible positions and vice versa. As a means of communication a black board agenda has been chosen to which each inference component signals new facts.

* From now on big letters will represent points and small ones name edges.




Fig. 16 The mapping from distance knowledge to position knowledge and vice versa. For each possible relation between the length of two edges of the triangle a , b , and c the possible positional relations within the reference frame are given. For the black dots the mapping can be converted meaningfully, i.e. one can map the position into a single relation between the lengths of the edges. For the gray dots every relation between the lengths of the edges are possible.

Note that the different logical combinations of the results of the mapping for each distance relation resemble the combination of the source relations. Thus, from $a < c$ and $b > c$ follows a sharper result because the intersection of the single results can be taken.




Fig. 17 The combination of more than one assertion. Note, that although within the qualitative spatial representation the shape of the restricted area and its small size can not be represented, this information is still available within the composed knowledge bases for means of visualization, for example.

The following Fig. 18 depicts the restrictions that are introduced by relating not only edges a , b , and c but also Dx , DyA , and DyB via first order \exists -calculus. The exact description of the areas and the corresponding constraints are not given due to the restricted publication space.




Fig. 18 The resulting areas from relating each of the edges to each other.

9 Conclusion

We have presented a framework for representing spatial knowledge and performing spatial reasoning. It features an intuitive iconic representation and is versatile. We have shown how the formalism can be used for spatial reasoning at different levels of granularity. The approach has been extended by the concept of motion and it has been combined with reasoning mechanisms for the domain of distances.

10 Bibliography

- Allen, J. F. (1983). "Maintaining Knowledge About Temporal Intervals." *CACM* **26**(11): 832-843.
- Cui, Z., A. G. Cohn and D. A. Randell (1993). *Qualitative Simulation Based on a Logical Formalism of Space and Time*. Proc. of the AISB Workshop on Qualitative and Causal Reasoning, Birmingham.
- Egenhofer, M. (1989). *A Formal Definition of Binary Topological Relationships*. W. Litwin and H.-J. Schek (Eds.): Proc. of the Third International Conference on Foundations of Data Organization and Algorithms, Springer, Berlin.
- Egenhofer, M. and R. Franzosa (1991). "Point-Set Topological Spatial Relations." *Intern. Journal of Geographical Information Systems* **5**(2): 161-174.
- Frank, A. U. (1991). *Qualitative Spatial Reasoning with Cardinal Directions*. Proc. of the Seventh Austrian Conference on Artificial Intelligence, Vienna, Springer, Berlin.
- Freksa, C. (1991). *Qualitative Spatial Reasoning*. D. M. Mark and A. U. Frank (Eds.): Proc. of the Cognitive and Linguistic Aspects of Geographic Space, Kluwer, Dordrecht.
- Freksa, C. (1992a). "Temporal Reasoning Based on Semi-Intervals." *Artificial Intelligence* **54**: 199-227.
- Freksa, C. (1992b). *Using Orientation Information for Qualitative Spatial Reasoning*. A. U. Frank, I. Campari and U. Formentini (Eds.): Proc. of the International Conference on Theories and Methods of Spatio-Temporal Reasoning in Geographic Space, Pisa, Italy, Springer, Berlin.

- Freksa, C. and K. Zimmermann (1992). *On the Utilization of Spatial Structures for Cognitively Plausible and Efficient Reasoning*. Proc. of the IEEE International Conference on Systems, Man, and Cybernetics, Chicago, IL.
- Güsgen, H. W. (1989). *Spatial Reasoning Based on Allen's Temporal Logic*. International Computer Science Institute, Berkeley, TR-89-049.
- Hernández, D. (1992). *Diagrammatical Aspects of Qualitative Representations of Space*. Proc. of the AAAI Spring Symposium on Reasoning with Diagrammatic Representations, Stanford.
- Hernández, D. (1992). *Qualitative Representation of Spatial Knowledge*. PhD Thesis, Institut für Informatik, Technische Universität München.
- Latecki L. and R. Röhrig (1993). *Orientation and Qualitative Angle for Spatial Reasoning*. To appear in Proc. of the IJCAI, Chambéry.
- Levitt, T. S., et al. (1987). *Qualitative Landmark-Based Path Planning and Following*. Proc. of the AAAI, Seattle, WA.
- Ligozat, G. (1993). *Qualitative Triangulation for Spatial Reasoning*. To appear in Proc. of the European Conference on Spatial Information Theory, Elba, Italy.
- Mukerjee, A. and G. Joe (1990). *A Qualitative Model for Space*. Proc. of the AAAI.
- Schlieder, C. (1990). *Akquisition räumlichen Wissens am Beispiel ebener Sicht- und Anordnungsverhältnisse*. W. Hoepfner (Ed.): Proc. of the Workshop Räumliche Alltagsumgebungen des Menschen, Koblenz-Landau, Universität Koblenz-Landau, Bericht 9/90.
- Zimmermann, K. (1993a). *Enhancing Qualitative Spatial Reasoning — Combining Orientation and Distance*. To appear in Proc. of the European Conference on Spatial Information Theory, Elba, Italy.
- Zimmermann, K. (1993b). *Measuring without Measures — The ?-Calculus*. Report.
- Zimmermann, K. and C. Freksa (1993). *Enhancing Spatial Reasoning by the Concept of Motion*. A. Sloman, D. Hogg, G. Humphreys, A. Ramsay and D. Partridge (Eds.): Proc. of the AISB, Birmingham, UK, IOS Press.