A Model of Spatial Position Based on Extremal Points

Raman Rajagopalan
Artificial Intelligence Laboratory
Department of Computer Science
University of Texas at Austin
Austin, Texas 78712
raman@cs.utexas.edu

Abstract

Two sets of fundamental properties which a spatial representation should describe are the position, extent, and orientation of objects, and topological and directional relations between objects. In previous approaches, the formalization of a particular property, such as topological spatial relations, has often been studied as an independent problem. The resulting theories have been elegant and compact, but may be difficult to use because of the assumptions on available knowledge. For example, several formalizations of topological relations assume knowledge of the boundaries of objects. Yet, knowledge of exact boundaries requires the knowledge of exact shapes, a property which is very difficult to describe symbolically.

Our approach was to first define an application, a magnetic fields problem solver capable of doing both spatial and dynamic reasoning, and a broad collection of qualitative problems to be solved before embarking on the task of defining a spatial representation. The nature of the solutions to the various problems provided clues on what spatial properties were of fundamental importance, and what information could be abstracted. We describe an adequacy criterion for abstractions of object shapes, and present a description of the spatial position and orientation of objects based on extremal points. We show that topological and directional relations may be readily described in terms of this object representation. Our spatial representation is compatible with existing tools for the qualitative simulation of dynamic systems, which allows us to reason about changes in spatial relations due to the translational and rotational motion of objects.

1 Introduction

Spatial reasoning has been a fundamental issue in several broad areas of Artificial Intelligence research including robot exploration and navigation, natural language processing, and in modeling the behavior of the physical world. Geographic information systems (GIS) must address issues in each of these subareas of research, and thus provide a good framework for the comparison, integration, and application of the many independently developed methods.

In robot exploration [Kuipers and Levitt 88], a fundamental goal is to learn a map of large scale space given only sensory inputs for the local space. A GIS may be provided with a map. However, both may share the goal of using an underlying map (spatial representation) for route planning and navigation.

Researchers interested in modeling the physical world have addressed the issues of motion and dynamic interactions between objects [Forbus, et. al 91]. A GIS may also have the need to model dynamic objects, such as a moving storm system. An automated navigation system for ships may compute routes using a static map, but may need to be aware of dynamic changes in the positions of storms.

Research in natural language (NL) processing [Retz-Schmidt 88] has addressed the issue of the communication of spatial information between agents. Goals include the generation of spatial descriptions, matching information in input descriptions against facts in an underlying knowledge base, and learning new facts from input descriptions. Important issues include deciding the applicability of a particular spatial preposition and deciding the frame of reference (intrinsic, deictic, extrinsic) to be used.

1.1 Existing Representations of Objects and Spatial Relations

A fundamental starting point for solving spatial reasoning problems is a representation which identifies the location, extent, and orientation of objects and spatial relations (topological and directional) between them. Several formalizations of topological relations (e.g., connected, equal, overlapped) are available [Abella and Kender 93, Cui, et. al 92, Egenhofer and Al-Taha 92, Hadzilacos and Tryfona 92]. Some of these include descriptions of how the topological relationship between two objects may change over time due to motion or changes in size or shape [Cui, et. al. 92, Egenhofer and Al-Taha 92]. Formalizations of directional relations, such as right-of, have been given in terms of global, intrinsic, and deictic frames of reference [Retz-Schmidt 88, Mukerjee and Joe 90].

Interestingly, although the formalizations are often qualitative, an actual reasoning system may require a fair amount of precise metric information to use the definitions. The definitions often assume that some particular piece of knowledge is available, when in fact, the information may have to be computed using numerically precise data. A GIS may contain vast amounts of metric information, but, a representa-

In: working notes of the Workshop on Advances in Geographic Information Systems, Arlington, Virginia, November 1993.

tion that gracefully degrades in performance in the presense of incomplete metric information will likely be superior to one that depends on very precise information.

Accidental dependencies on very precise metric information can defeat one of the purposes of a qualitative formalism, which is to allow reasoning about a physical situation even when absolute knowledge of the values of variables is not available. However, the existing formalizations are still a good starting point for further work, since they succeed in mapping many physical situations into a much smaller set of qualitative situations. This facilitates reasoning about the most important issues and allows one to abstract away unimportant details.

For example, Cui, Cohn, and Randell [1992] and Hadzilacos and Tryfona [1992] assume knowledge of the boundaries and interiors of objects, and define topological relations with respect these properties. However, exact shapes, and thus, exact boundaries, are very difficult to describe. In natural language, there are relatively few symbolic terms to describe specific shapes. In human communication, an abstraction term, such as 'polygon' or 'closed object' may be used to partially describe a complex shape, but a drawing may be required to convey the exact shape.

In describing the CITYTOUR system, Retz-Schmidt [1988] describes an algorithm for reasoning about directional prepositions with intrinsic and deictic frames of reference. Right, left, front, and back half-planes are to be drawn based on the intrinsic front or accidental front (in deictic use) of the reference object. The directional relation between the primary object and the reference object is given by the half-plane in which the primary object lies. Although left unspecified in Retz-Schmidt's survey article, we assume that the half-planes are computed as required using precise metric information.

Another example is the collision parallelogram representation of Mukerjee and Joe [1990]. Projections called lines of travel are drawn parallel to the surface normal direction of the intrinsic front of an object. When the lines of travel of two objects intersect, they form a parallelogram. Mukerjee and Joe provide a rich theory of how directional relations between two objects may be derived given knowledge of the positions of objects and the collision parallelograms. The collision parallelograms may have to be computed and maintained using precise metric information, particularly if objects are allowed to rotate.

1.2 Abstractions of Shape

Spatial representations often make use of abstractions to hide unimportant details and to simply the reasoning process. A fundamental spatial property for which abstractions are used is the shape (spatial extent) of an object. Rectangular bounding boxes are a commonly used abstraction method, but the criteria for drawing the boxes may vary. Retz-Schmidt [1988] and Mukerjee and Joe [1990] draw bounding boxes around the 'front' of an object. Abella and Kender [1993] draw a bounding box based on a mathematical abstraction, inertia tensors. Cui, Cohn, and Randell [1992] use the convex hull as the abstraction method.

We feel that the choice of the abstraction method for shapes is fundamental. The chosen method can affect the way in which spatial relations are expressed, which in turn may affect the reasoning power of the application system. For example, Kuipers [Kuipers and Levitt 88] represents distinctive places as points and connections between them as lines. While this approach is quite appropriate for navigation and route finding, queries about the intrinsic fronts of objects are inexpressible, since objects are represented as points.

- **Diagram:**

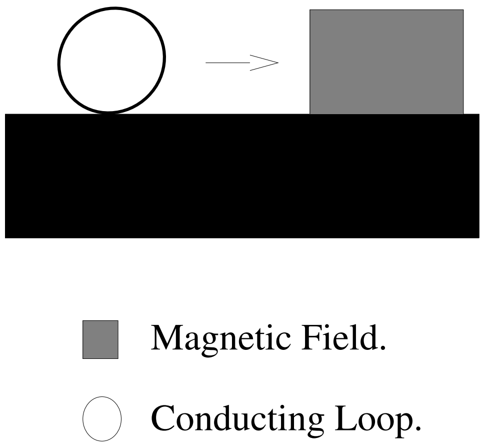

- **Text:** The loop is rolling along the block.

Figure 1: Example of text and diagram input for a magnetic fields problem.

In the next section, we describe a representation for two and three dimensional objects based on the extremal points of objects. This representation is both intuitive and powerful. Starting with our simple representation of object position and extent, we can easily determine the orientation of the object. Further, we show that many topological and directional spatial relations may be defined in terms of inequalities between the extremal points of different objects. This representation allows reasoning about changes in relative positions due to the translational motion of objects.

Our spatial representation readily fits existing methods for the qualitative simulation of dynamic systems. We have implemented descriptions of spatial relations as model fragments for the qualitative process compiler, or QPC [Crawford, et. al. 90], which in turn uses the QSIM qualitative simulation system [Kuipers 86]. We also include a library of processes which can modify the spatial state, such as the motion of objects. QPC performs of the task of determining the current spatial state, determines what dynamic changes are taking place, and detects significant changes in the qualitative spatial state as they occur. We are using these models in a problem solving system which reasons about the effects of magnetic fields [Rajagopalan and Kuipers 93a]. Our implementation architecture will also permit us to inherit any advances in the underlying systems, QPC and QSIM, such as the ability to incorporate incomplete quantitative information into the models [Kay and Kuipers 93].

2 A Spatial Representation Based on Extremal Points

Figures 1 and 2 illustrate two of the magnetic fields problems we address. In the first, there is a conducting object which is rolling along a block into a magnetic field. As in textbooks, we input the initial spatial configuration through a diagram and provide additional text to describe any dynamic changes and other problem specific information. In the diagram, we use patterns and shading to denote par-

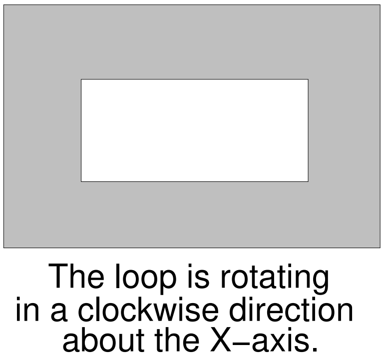

Figure 2: A rudimentary generator: A conducting loop rotating in a magnetic field.

ticular types of domain objects. This additional domain knowledge provides the means to match object descriptions in the diagram to object references in text. The methods for integrating the information in the diagram and text input into a single knowledge base are described separately [Rajagopalan and Kuipers 93b].

For the problem in Figure 1, as the conducting loop enters and exits the field, the area of overlap between the two will change. This causes an induced emf which in turn causes a current flow in the loop. Figure 2 shows a loop of wire inside a magnetic field. An external agent is turning the loop. As the loop rotates, the projected area of overlap between the loop and the field will change. This again results in a current flow in the loop.

Both problems require reasoning about changes in area of overlap between a conducting object and a magnetic field. In the first case, the change in the area of overlap is due to translational motion. In the latter, the change is due to rotational motion. These aspects of the problem involve reasoning about *relative positions*. Once we establish a current flow in the loop, we have to use *orientation* information to determine whether a force will act on the loop. If so, we will need to determine where the force acts, decide if there is an effect on the translational velocity and whether the force exerts a torque.

In these problems, we are given a qualitative description of the initial state that may map to many physical instances. We are to reason about the consequences of qualitatively significant changes in the relative positions of 2-dimensional objects. The exact shape of objects is not of much concern, and an approximation of the region occupied by an object may be used. However, to be able to answer such questions as 'Will a given force cause a torque?', we will need more precise information about the orientation of an individual object.

We use two abstractions to approximate the spatial extent of a convex two dimensional object, as illustrated in Figure 3. When only translational motion is possible, we use a **rectangular bounding box** drawn around the object's **extremal points** - the `rightmost` (rm) , `leftmost` (lm), `topmost` (tm), and `bottommost` (bm) points. When rotational motion is also possible, we use a **bounding circle**, drawn to trace out all the points covered by the object if it is allowed to rotate about its center of gravity. The center of the circle is the center of gravity, and the radius is given by the distance from the center of gravity to the furthest point on the perimeter.

We treat the extremal points of the bounding circle as

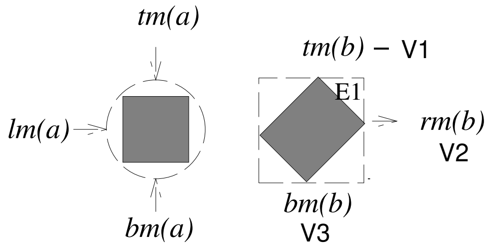

Figure 3: Examples of bounding boxes and bounding circles as abstractions of shape for two polygons.

the extremal points of the enclosed object. This ensures that the enclosed object is always fully contained within the region of space given by the abstraction used. For three dimensional objects, bounding cubes and bounding spheres can be defined in a similar manner using two additional extremal points - the frontmost and rearmost points. While other abstractions of shape, such as the convex hull [Cui, et. al. 92] may provide tighter bounds, they may require numerically precise data to compute and maintain. Our abstraction methods have the advantage that they can be maintained and used with very little information (four or six points for each two or three dimensional object, respectively) even when objects may translate and rotate.

The use of the bounding circle approximation ensures that rotation about the center of gravity cannot change the spatial extent an object, since the actual object always remains within the bounding circle. Note however, that the orientation of the object may change with rotation, and will have to be modeled separately. To reason about objects in translational motion, we simply need to note that the topmost and bottommost values of the bounding box or bounding circle will change when the Y-velocity of the object in non-zero, and that the rightmost and leftmost values will change when the X-velocity is non-zero.

Thus, by defining spatial relations through inequality relations between the extremal points of objects, we can easily determine qualitatively significant changes as they occur. For example, the inequality relation $rightmost(A) < leftmost(B)$, captures the fact that object A is to the left of B. Then, if object A is in motion, and has a positive X-velocity, we may not only conclude that the rightmost(A) will increase, we may also use the inequality relation to recognize that a qualitatively significant change will occur when rightmost(A) equals leftmost(B).

2.1 An adequacy criterion for abstractions of shape

An abstraction of shape is an incomplete description of the actual extent of an object. It is possible for the abstractions of two objects to satisfy a spatial property, such as overlap, even though there may be physical situations when the actual positions of the objects do not satisfy the property. In our work, which involves reasoning about possible interactions and overlaps among extended objects in 2-dimensional and 3-dimensional space, we have found it to be useful to select an abstraction method that at least allows the conclusion that two objects *do not* satisfy a spatial relation if the abstractions do not satisfy the relation. We have adopted the following adequacy criterion to evaluate potential abstractions for our purposes:

- Any abstraction for the spatial extent of an object should fully cover the actual extent of the object at all times.

Overlooking this adequacy criterion may lead to counterintuitive conclusions. For example, Abella and Kender [1993] define 'near' to be true when the abstractions of shape (bounding boxes) of two objects intersect. However, visual inspection of their figures reveals that their method of drawing the bounding box does not completely cover the object, and that it is possible for two objects to actually intersect, and not be considered 'near'.

An abstraction of spatial extent is useful only if it is easier to maintain than knowledge of the actual position and orientation of an object. This observation, along with the adequacy criterion, led to our use of two abstractions, the rectangular bounding box and the bounding circle, instead of one.

Consider the two rectangles in Figure 3. Allow the right object to rotate. If the bounding box abstraction is used, we not only have to decide how the extremal points change, we also have to decide if the rotation can cause a change in the spatial relationship with the left object. If the two objects are close enough, the two rectangles may intersect. Without precise information about the distances between the objects and the sizes of the objects, we cannot make this decision. The bounding circle abstraction ensures that rotation cannot change the spatial relation with another object. The penalty for using the bounding circle approximation is that it is more *incomplete*. Notice that the bounding circles intersect in Figure 3 although the bounding boxes do not. The use of the bounding box abstraction would allow us to conclude that the two rectangles are *disconnected*. We would conclude that they are *overlapped* if bounding circles are used.

2.2 Modeling Relative Positions and the Effects of Translational Motion

Below, we present definitions for several spatial relations used in reasoning about the example in Figure 1. These definitions are implemented as model fragments for the qualitative process compiler (QPC) [Crawford, et. al. 90]. QPC model libraries may be constructed in a modular fashion. Therefore, new spatial relations may be defined and included in the library as required. We have used such formalizations of those of Cui, Cohn, and Randell [1992] to build our library of spatial relations.

Definitions for the spatial relations left-of, right-of, aligned-x, aligned-y, overlapped, enclosed, overlap-left-boundary, and overlap-right-boundary, are given below. Aligned-X(A,B) is true if the objects can intersect through motion in the X direction, and Aligned-Y(A,B) is true if the two objects can intersect through motion in the Y-direction. The terms rm, lm, tm, and bm refer to the rightmost, leftmost, topmost, and bottommost points respectively. The terms rightmost and leftmost are X-coordinate values and the terms topmost and bottommost are Y-coordinate values. Actual coordinate values are not used. We only maintain inequality relationships between pairs of points.

- `Left-of(A,B)`: $rm(B) < lm(A)$.
- `Aligned-X(A,B)`: $tm(A) > bm(B) \wedge bm(A) < tm(B)$.
- `Aligned-Y(A,B)`: $rm(A) > lm(B) \wedge lm(A) < rm(B)$.
- `Overlapped(A,B)`: $Aligned{-}X(A,B) \wedge Aligned{-}Y(A,B)$.
- `Overlap-left-boundary(A,B)`: $Overlapped(A,B) \wedge rm(A) \leq rm(B) \wedge lm(A) < lm(B)$.
- `Enclosed(A,B)`: $Overlapped(A,B) \wedge rm(A) \leq rm(B) \wedge lm(A) \geq lm(B) \wedge tm(A) \leq tm(B) \wedge bm(A) \geq bm(B)$.
- `Overlap-right-boundary(A,B)`: $Overlapped(A,B) \wedge lm(B) \leq lm(A) \wedge rm(A) > rm(B)$.
- `Right-of(A,B)`: $rm(B) < lm(A)$.

In figure 1, the initial state may be defined as follows. We ignore the block as it has no bearing on the solution to the problem. Note that in the initial state, two spatial relations are true: Left-of(loop, field) and Aligned-X(loop, field).

1. lm(loop) < lm(field), lm(loop) < rm(field).
2. rm(loop) < lm(field), rm(loop) < rm(field).
3. tm(loop) < tm(field), bm(loop) < tm(field).
4. tm(loop) > bm(field), bm(loop) = bm(field).
5. x-velocity(loop) > 0.

By defining spatial relations in terms of inequality relations between variables, we can examine the direction of change in the values of the variables to determine if a spatial relationship might change. For example, in the initial state, Aligned-Y is not true since rm(loop) > lm(field) is not true. However, the knowledge that rm(loop) is increasing (X-velocity is positive) can be used to determine that the relation will eventually be satisfied. Once this is satisfied, so are the conditions for Overlapped(loop,field) and Overlapped-left-boundary(loop, field). As the loop proceeds to the right, it will become enclosed in the field, and then exit to the right of the field. Eventually, the loop will be to the right of the field.

We are building a magnetic fields problem solver as an application for our spatial reasoning methods. The problem solver uses domain models to recognize spatial configurations of significance in the magnetic fields domain and to determine their dynamic effects. In Figure 4, we show the QPC output for the example in Figure 1. The methods for reasoning about relative positions and translational motion have been successfully implemented, and we are currently implementing the portions of the domain theory that require knowledge of object orientation, such as methods for describing the points of action of magnetic forces and the dynamic effects of these forces.

In Figure 4, the loop starts to the left of the field, and is moving to the right. The values of the rightmost and leftmost points of the loop are steadily increasing. After some time, the loop begins to move into the field, and the area of overlap is positive and increasing. Once the loop is fully enclosed in the field, the area of overlap remains constant until the loop begins to exit the field to the right. The area of overlap then decreases to zero, and becomes undefined once the loop completely exits the field. When the area of overlap between the loop and the field is changing, so does the flux through the loop. This results in an induced

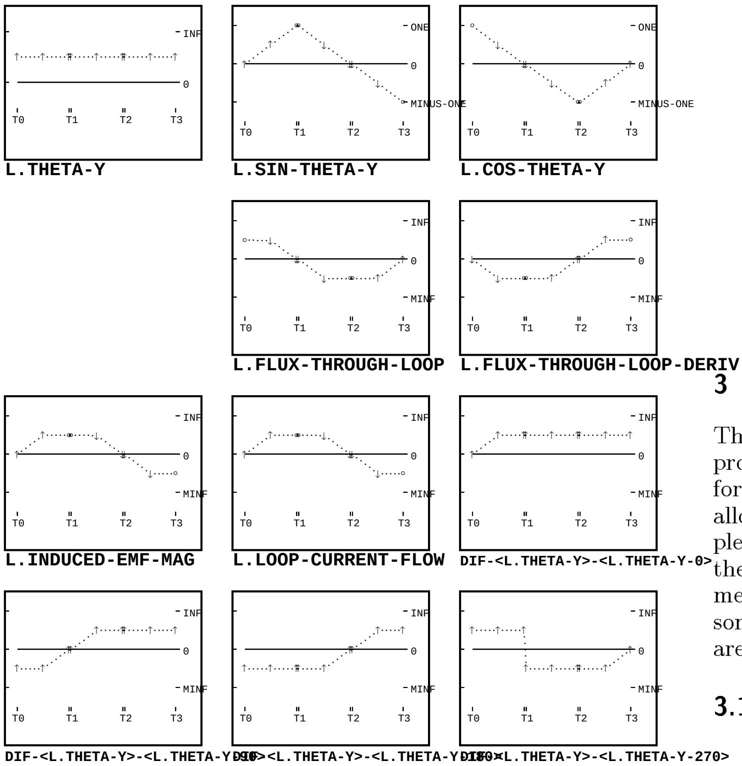

Figure 4: Qualitative behavior produced by QPC for the example in Figure 1.

emf being set up in the loop which in turn causes a current flow in the loop.

The variables whose names begin with 'DIF' are QPC generated variables which encode the inequality relationships of interest in the simulation. Notice that the qualitatively interesting changes occur whenever one of the 'DIF' terms crosses zero. For example, when the variable DIF-<F.LM-XVAL>-<L.RM-XVAL> is positive, the loop is to the left of the field. The value becomes zero at the instant when the loop and the field meet.

2.3 Modeling the Orientation of Objects

For convex objects, the representation of position based on extremal points also gives the orientation with respect to a global Cartesian frame of reference. We define the **orientation** of an object as a union of pairs (s_i, d_i), denoting some surface of the object and its surface normal direction. The set of all s_i should exactly cover the surface area of the object.

The directions, d_i, are defined by a qualitative vector (Xmag, Ymag, Zmag) where each of Xmag, Ymag, and Zmag may carry a value from the set $\{-, 0, +\}$ to denote whether the direction vector has a component along each of the X, Y, and Z axes respectively.

Given this definition of orientation, we may determine the orientation of any surface of a two-dimensional convex object in the X-Y plane as follows:

- For a polygon, if an edge forms an extremal point, then the orientation of that edge will be (1 0 0) for the rightmost point, (0 1 0) for the topmost point, (-1 0 0) for the leftmost point, and (0 -1 0) for the bottommost point.
- For any surface which lies between the topmost and rightmost points, the orientation vector is (1 1 0).
- For any surface which lies between the rightmost and bottommost points, the orientation vector is (1 -1 0).
- For any surface which lies between the bottommost and leftmost points, the orientation vector is (-1 -1 0).
- For any surface which lies between the leftmost and topmost points, the orientation vector is (-1 1 0).

3 Future Work

There are two directions for future work in this ongoing project. The immediate problem is to implement a method for modeling the orientation of a rotating object. This will allow us to address such problems as the generator example given in Figure 2, which describes a loop rotating about the X-axis. Another topic for further research is to explore methods for applying our results to problems involving reasoning about intrinsic and deictic frames of reference, which are of great importance in natural language understanding.

3.1 Determining the Orientation of Rotating Objects

In Section 2, we introduced the bounding circle abstraction for modeling the spatial extent of objects which may rotate. This abstraction was shown to be useful for reasoning about changes in relative position. As an object rotates, the identities of the vertices which form the extremal points will change. To use the methods for representing orientation described in the previous section, we must explicitly model the changes in the identities of the extremal points as an object rotates.

We have discovered a promising solution to the problem, but the methods have yet to be tested via implementation. Therefore, we only provide a sketch of the issues here. For reasoning about rotation about the X and Y axis, we maintain an understanding of the angle between the plane of the rotating (two-dimensional) object and the X-Y plane. Knowledge of this angle and its direction of change will allow us to determine if the projected area of the object onto the X-Y plane is increasing or decreasing. For the generator problem of Figure 2, this will allow us to determine if the flux through the loop is increasing or decreasing as it rotates. We can also use knowledge of the angle between the plane of the object and the X-Y plane to determine if the topmost and bottommost points are reversed (rotation about X axis), or if the rightmost and leftmost points are reversed (rotation about Y axis).

Problems involving rotation about the Z-axis are more complicated. Here, we will use knowledge of the internal angles of a polygon to determine the sequence in which the identities of the extremal points will change. Given that vertex V2 is the topmost point of an object rotating in the clockwise direction, and that vertex V1 follows V2, we may conclude that V1 will be the next topmost point. The difficult problem is to determine which extremal point will change first. Given that V3 is the rightmost point and that V2 follows V3, the question now becomes will V1 becomes the topmost point before V2 becomes the rightmost point? Since the bottommost and leftmost points may also change, there is always a four-way race condition.

Given that θ is the internal angle of some vertex V, once V becomes an extremal point, it will remain an extremal point over a total rotation of $180 - \theta$ degrees before it is replaced by some other vertex. This observation is the key behind our proposed solution. We can maintain an ordered list, ($\theta 1$, $\theta 2$, $\theta 3$, $\theta 4$), which stores the number of degrees an object must be rotated for the identity of each of the four extremal points to change. A qualitatively interesting change will occur only after a rotation of $\theta 1$ degrees. Then, we may subtract $\theta 1$ from each of the other values and remove $\theta 1$ from the list. Assume that vertex V is the new extremal point and that θ_V is its internal angle. We must then insert the angle (180 - θ_V) into the ordered list. Knowledge of connectivity and the internal angles will be required to use this solution. We are also investigating a qualitative method which does not require precise knowledge of the values of the internal angles.

4 Intrinsic and Deictic Frames of Reference

We discussed in Section 1 that methods for spatial reasoning are being studied by researchers in several distinct subareas of Artificial Intelligence. Our work fits in the category 'methods for reasoning about the physical world.' Problems in this category typically require a global frame of reference. Researchers in natural language understanding, on the other hand, have to consider the possibility that a spatial relation could be described with respect to intrinsic and deictic frames of reference as well as extrinsic frames of reference.

Mukerjee and Joe [1990] and Retz-Schmidt [1988] both present methods for reasoning about intrinsic frames of reference. The methods involving drawing a rectangular bounding box around the intrinsic front on an object and determining the left, right, front, and back directions with respect to this intrinsic front.

In our work, we provide methods for determining the orientation, or surface normal direction, of any surface of any object. We can also obtain the relative position of any two objects with respect to a global frame of reference. We are now addressing the question 'Is there a mapping between an intrinsic frame of reference to a global frame of reference which will allow the relative positions of two objects to be compared with respect to the intrinsic frame of reference, assuming that only qualitative information is available?'.

5 Acknowledgements

I wish to thank my dissertation advisor, Professor Benjamin Kuipers, for providing many useful suggestions to guide this research and James Lester for reading early drafts of this paper and providing valuable comments. This work has taken place in the Qualitative Reasoning Group at the Artificial Intelligence Laboratory, The University of Texas at Austin. Research of the Qualitative Reasoning Group is supported in part by NSF grants IRI-8905494 and IRI-8904454, by NASA grant NAG 2-507, by the Texas Advanced Research Program under grant no. 003658175, and by the Jet Propulsion Laboratory.

6 References

[Abella and Kender 93] Abella, A. and J. Kender. Qualitatively Describing Objects Using Spatial Prepositions. In: AAAI-93, Washington, D.C., 1993.

[Crawford, et. al. 90] Crawford, J., A. Farquhar, and B. Kuipers. QPC : A Compiler from Physical Models into Qualitative Differential Equations. In: AAAI-90, Boston, MA, 1990.

[Cui, et. al. 92] Cui, Z., A. Cohn, and D. Randell. Qualitative Simulation Based on a Logical Formalism of Space and Time. In: AAAI-92, San Jose, CA, 1992.

[Egenhofer and Al-Taha 92] Egenhofer, M. and K. Al-Taha. Reasoning about Gradual Changes in Topological Relationships. In: A. Frank, I. Campari, and U. Formentini, Eds., *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space*, pp. 196-219, Springer-Verlag, Berlin, 1992.

[Forbus, et. al. 91] Forbus, K., P. Nielsen, and B. Faltings. Qualitative Spatial Reasoning: the CLOCK Project. *Artificial Intelligence* 51, 417-471, 1991.

[Hadzilacos and Tryfona 92] Hadzilacos, T., and N. Tryfona. A Model for Expressing Topological Integrity Constraints in Geographic Databases. In: A. Frank, I. Campari, and U. Formentini, Eds., *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space*, pp. 252-268, Springer-Verlag, Berlin, 1992.

[Kay and Kuipers 93] Kay, H., and B. Kuipers. Numerical Behavior Envelopes for Qualitative Models. In: AAAI-93, Washington, DC, 1993.

[Kuipers 86] Kuipers, B. Qualitative Simulation. *Artificial Intelligence* 29, 289-338, 1986.

[Kuipers and Levitt 88] Kuipers B., and T. Levitt. Navigation and Mapping in Large-Scale Space. AI Magazine, 9(2), 25-43.

[Mukerjee and Joe 90] Mukerjee, A. and G. Joe. A Qualitative Model for Space. In: AAAI-90, Boston, MA, 1990.

[Rajagopalan and Kuipers 93a] Rajagopalan R., and B. Kuipers. Qualitative Spatial Reasoning about Objects in Motion: Application to Physics Problem Solving. Manuscript submitted to CAIA-94, San Antonio, Texas.

[Rajagopalan and Kuipers 93b] Rajagopalan R., and B. Kuipers. The Figure Understander: A System for Integrating Diagram and Text Input to a Knowledge Base. Manuscript to be submitted to: IEA/AIE-94, Austin, Texas.

[Retz-Schmidt 88] Retz-Schmidt, G. Various Views on Spatial Prepositions. *AI magazine*, 9(2), 95-105.