In *Working Notes of the AAAI Workshop on Spatial and Temporal Reasoning*
AAAI-94, Seattle, Washington, 1994.

On Shape Abstractions for Qualitative Spatial Reasoning *

Raman Rajagopalan
Department of Computer Sciences
University of Texas at Austin
Austin, Texas, 78712

Abstract

The representation of the shape of an object is the foundation of any theory of spatial reasoning. Many spatial representations have been developed which use simplifying abstractions to approximate the actual shapes of objects. A limitation of these methods is that the abstractions themselves are mathematical entities, making them less useful for solving problems where only qualitative descriptions of the initial state are available. We discuss the advantages and limitations of using shape abstractions and present our spatial representation, which, unlike the previously developed methods, requires only qualitative information to describe and maintain, even in dynamically changing worlds.

1 Introduction

Spatial reasoning has been studied in Artificial Intelligence from a number of perspectives, including natural language understanding [Freksa 92, Maass 94, Mukerjee and Joe 90, Retz-Schmidt 88], robot exploration and mapping [Kuipers and Levitt 88], and qualitative reasoning about physical systems [Cui, Cohn, and Randell 92, Forbus, Nielsen, and Faltings 91, Joskowicz and Sacks 91, Rajagopalan 94]. A common starting point is the representation of the spatial extent of an object.

Herskovits [1985] and Landau and Jackendoff [1993] have documented the use of shape abstractions in Human reasoning about spatial relations. Herskovits provides examples involving the use of several different shape abstraction methods, including points, boundary outlines, and volumes. Several theories have since been developed for qualitative spatial reasoning using simplying shape abstractions. These theories assume that the abstractions are provided as part of the input description, and do not provide a means to *qualitatively* model the abstractions. They also do not discuss how to compute changes in the position or orientation of the abstractions, or the objects within, as objects translate and rotate. As many of the shape abstractions are *quantitative* entities, they are unsuitable for solving problems where only qualitative information is available about the initial spatial state.

In this paper, we discuss the properties of the spatial representations that have previously been used for qualitative spatial reasoning (Section 2), present the requirements of our application, a magnetic fields problem solver (Section 3), and discuss the advantages and limitations of using shape abstractions (Section 4). We then describe our spatial representation (Section 5) and show that unlike many of the previously developed methods, it can be described and maintained with only qualitative information, even in dynamically changing worlds.

2 Representations in Related Work

Many spatial representations have been proposed for qualitative spatial reasoning. Problems which require reasoning about exact shapes, such as the mechanical constraint problems addressed by Forbus, Nielsen, and Faltings [1991] and Joskowicz and Sacks [1991], have inspired the development of qualitative theories for reasoning with such complex quantitative spatial representations as the configuration space. When knowledge of exact shapes is unnecessary, simplifying shape abstractions have been used. For example, Kuipers [Kuipers and Levitt 88] uses a point-based representation to describe the position of objects in a topological map.

In our work, we consider problems which require reasoning about the properties of two and three dimensional objects, such as intersection and containment. In previous work, two and three dimensional spatial representations have been developed based on the use of convex hulls [Cui, Cohn, and Randell 92], inertia tensor-based bounding rectangles, [Abella and Kender 93], bounding rectangles based on intrinsic front directions [Retz-Schmidt 88, Mukerjee and Joe 90], and cuboids [Maass 94] as shape abstraction methods.

These theories do not provide a qualitative model for the shape abstractions, and instead assume that they are provided as part of the input description. Although the details are left unspecified, we can only assume that

*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-8904454, IRI-9017047, and IRI-9216584, and by NASA contracts NCC 2-760 and NAG 9-665.

the implementations of these theories make use of quantitative knowledge of the spatial state.

For example, the inertia tensor-based rectangular bounding box used by Abella and Kender and the convex hull used by Cui, Cohn, and Randell are both quantitative entities. Thus, the qualitative models of spatial relations developed based on these abstractions cannot be implemented without quantitative knowledge of the spatial state. Abella and Kender define two objects to be near each other if their bounding boxes intersect. Similarly, Cui, Cohn, and Randell define two objects to be intersecting if their convex hulls intersect. In both cases, quantitative knowledge will be required to determine if an intersection exists.

Retz-Schmidt [1988] and Mukerjee and Joe [1990] address problems in Natural Language Understanding involving the need to reason with arbitrarily oriented frames of reference. Freksa [1992] shows that, in general, the problem is underconstrained if only qualitative information is available about the positions of objects. The CITYTOUR project [Retz-Schmidt 88] uses rectangular bounding boxes drawn around the intrinsic fronts of objects to define forward, backward, left, and right half-planes with respect to a given object. Mukerjee and Joe [1990] define a rectangular bounding box in a similar manner, but compute a *collision parallelogram* to reason about the relative positions of two objects.

In both cases, it appears that the abstractions are computed precisely based on knowledge of the shapes of the underlying objects. Otherwise, from the results of Freksa [1992], one would expect much more ambiguity in placing objects within half-planes, or in computing the locations of the collision parallelograms. Further, the theories do not specify how the regions of the half-planes or the collision parallelograms can be recomputed with only qualitative information if the underlying objects are allowed to rotate.

3 The Magnetic Fields Problem Solver

In our work, we address qualitative problems from the magnetic fields domain [Resnick and Halliday 88] whose solutions require integrated spatial and dynamic reasoning about dynamically changing worlds where objects are free to translate and rotate. Our problems cover the basic principles of operation of such commonly used devices as motors, generators, and transformers.

Two representative problems are shown in Figures 1 and 2. The interested reader may consult [Rajagopalan 94] for further details on the implementation of the problem solver, and its solution for these problems. Here, we only wish to illustrate the requirements for spatial reasoning. Our problems require reasoning about the effects of changing the magnetic flux passing through a closed conducting object. The change in flux can be due to a time-varying magnetic field (transformer), or due to the translational (Figure 1) or rotational (generator, Figure 2) motion of the conducting object with respect to the field.

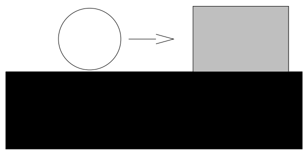

Figure 1: A magnetic fields problem involving translational motion: What happens as the circular conducting loop (white) rolls into the magnetic field (gray)?

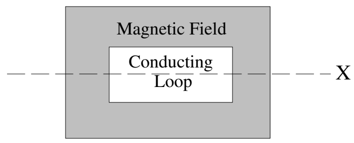

Figure 2: A magnetic fields problem involving rotational motion: A rudimentary generator. What happens as the conducting loop (white) rotates in a counterclockwise manner about the X-axis?

The most interesting spatial relations involve intersection between a closed conducting object and a magnetic field. For problems involving motion, we must infer any changes in the area of overlap, and therefore the magnetic flux passing through the conducting object, and determine if an electromagnetic force (emf) could be induced in the conducting object.

We require a model of the orientation of objects to represent the directions of magnetic fields and the current flow in any surface of a conducting object. This information is used to compute the line of action and the dynamic effects of any magnetic forces that may act on the conducting object.

4 Properties of Shape Abstractions

Qualitative simulation systems such as QSIM [Kuipers 86] provide the soundness guarantee that all the real behaviors that can follow from a given initial state will be represented in the output of the simulation. When an abstraction is used to *approximate* the actual spatial extent of an object, there can exist real behaviors that can no longer be reasoned about. Therefore, spatial reasoning systems should provide an explicit statement about the guarantees that they provide.

If a maximally covering abstraction is used, such as the bounding rectangle used in Figure 3, one can *guarantee* non-intersection conclusions based on the abstractions to be sound, but not intersection conclusions (ex-

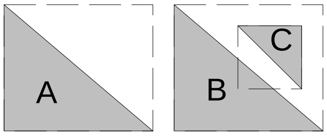

Figure 3: Example of a limitation of using shape abstractions: Although triangles B and C are disjoint, one could falsely conclude that object C is inside object B based on the relationship between their rectangular shape abstractions.

cept in special cases). Similarly, if an abstract shape is fully contained within the actual shape, then intersection conclusions can be guaranteed to be sound, but not non-intersection conclusions.

For example, in Figure 3, since the bounding rectangle for object A does not intersect any other object, we may soundly conclude that object A cannot be intersecting any other object. However, it's possible that the underlying objects do not intersect even if the maximally covering rectangles intersect. In Figure 3, triangles B and C are not intersecting, but the bounding rectangle of B encloses that of C.

Our problems require reasoning about what can happen *if* a conducting object partially overlaps or is fully enclosed within a magnetic field. We therefore enforce the constraint that a shape abstraction should fully enclose an underlying object at all times to guarantee when it is at all possible for two objects to be intersecting. We prove weaker conditions, such as the fact that the bounding shapes of two objects are intersecting, and then make use of explicit assumptions to consider the possibility that the underlying physical objects satisfy the same property.

5 A Qualitative Spatial Representation Based on Extremal Points

We employ a qualitative spatial representation which requires knowledge of only the connectivity of an object and the identities of its **extremal points**: the `topmost` (tm), `rightmost` (rm), `bottommost` (bm) `leftmost` (lm), `frontmost` (fm), and `rearmost` (rrm) points. We describe and maintain knowledge of both directional and topological **spatial relations** between objects based on *inequality relations* between their extremal points. We directly infer the **orientation** of any surface of an object with respect to a global Cartesian frame of reference from knowledge of the *identities* of the object's extremal points.

5.1 Modeling Spatial Extent

We use two abstractions, **a rectangular bounding box** drawn around the extremal points of an object, and a **bounding circle**, as described below, to approximate the region occupied by the object. For problems in three-dimensional space, we use a **bounding cube** and a **bounding sphere** respectively. In Figure 4, we illustrate our abstraction method on two polygons. We draw a bounding circle around object A and a rectangular bounding box around object B. The radius of the bounding circle is the maximum distance from the center of gravity of an object to any point on its perimeter.

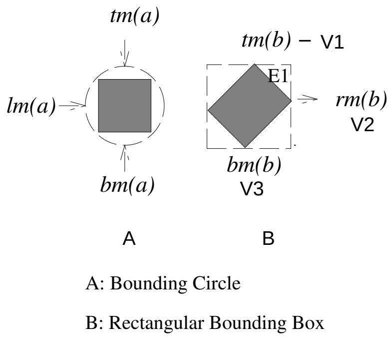

Figure 4: Examples of extremal point-based bounding rectangles and bounding circles as methods of shape abstraction for two polygons.

The bounding box abstraction is sufficient for reasoning about problems involving static worlds or only translational motion. The bounding circle/sphere abstraction is used for problems involving rotation about the center of gravity. Since the identities of the extremal points of an object can change during rotation, the shape of a bounding box may not be constant during rotation of the underlying object. The bounding circle abstraction will not only retain a fixed shape, it also provides the guarantee that a rotating object will always remain within the circle.

As discussed in Section 4, maximally covering abstractions can be used to soundly infer that two objects are not intersecting, but it is possible that the abstract shapes intersect even when the underlying objects don't. The bounding circle abstraction has the additional limitation that it cannot be used by itself to reason about changes in intersections due to rotation. If two bounding spheres are not intersecting, one can soundly conclude that the underlying objects cannot come into contact after any rotation of either object about its center of gravity. However, when it is possible for two objects to come into contact through the rotation of either object, or for an existing contact to be removed, the bounding circles of the objects will *always* be intersecting. Additional information, such as the distances between the objects, their sizes, and relative orientations, will be required to determine with certainty if rotation can create or remove contacts between the underlying objects.

5.2 Relative Position and Orientation

We describe the relative positions of objects through *inequality relations* between their extremal points. If the rectangular/cubic bounding box abstraction is used, the extremal points of the shape abstraction and the underlying object are the same. If the bounding circle/sphere abstraction is used, the extremal points used are those of the bounding abstraction, and not the underlying object. For example, in Figure 1, the inequality relation $rm(loop) < lm(field)$ models the spatial relation that the conducting loop is to the left of the magnetic field.

In general, for relations involving intersections between objects, we can prove only the weaker relation that the bounding abstractions of the objects satisfy the desired property. For example, the field (f) `bounding-abstraction-encloses` the loop (l) if $(lm(f) \leq lm(l)) \wedge (rm(f) \geq rm(l)) \wedge tm(f) \geq tm(l)) \wedge (bm(f) \leq bm(l))$.

We define the `orientation` of a surface of an object through its surface normal direction [Nielsen 88]. In our representation, this may be obtained through knowledge of the *identities* of the extremal points between which a surface lies. For example, the orientation vector of edge E1 for object B in Figure 4 has positive X and Y components since it lies between Vertices V1 and V2, the topmost and rightmost points. This model is sufficient for inferring the direction of current flow in a conductor, which is tangent to the surface, or the direction of any magnetic force which acts on the conductor, which will be normal to the surface.

Rectangular bounding boxes and bounding circles are powerful abstractions for qualitative spatial reasoning: numerically precise information is not required, even in a dynamically changing world, to maintain the bounding boxes and bounding circles. All we need are the inequality relationships between the extremal points of the bounding abstractions that are used. The use of circular/spherical abstractions for rotating objects ensures that only translational motion can change the position of the extremal points of the shape abstractions and thus, the relative positions of objects (assuming that objects have fixed shapes).

Orientation information may be obtained if the connectivity of an object and its extremal points are known. Only rotational motion can change the identities of the vertices and edges that form the extremal points of an object, and thus, its orientation.

5.3 Translational Motion

Our inequality relation-based model of relative positions allows us to recognize that a qualitatively interesting event will occur when the inequality relation changes (e.g., when $rm(loop) = lm(field)$ in Figure 1). This can be used to study the effects of translational motion on the spatial state.

The values of the rightmost and leftmost points of an object will change over time if the X-velocity of the object is non-zero, and the values of the topmost and bottommost points will be affected by the Y-velocity. In Figure 1, the X-velocity of the loop is positive, and we recognize that the distance between the rightmost point of the loop and the leftmost point of the field will be decreasing. We then infer that a qualitatively interesting state change will occur once $rm(loop) = lm(field)$.

5.4 Rotational Motion

With the use of the bounding circle/sphere abstraction, rotation can only change the orientation of an object. To maintain the orientation of a rotating object, we must determine the changes in the identities of the vertices and edges that form its extremal points as it rotates. Currently, we only consider the rotation of two dimensional objects about an axis that passes through their center of gravity. Rotation can also change the projected area of an object onto another object in a different plane, a property of great interest for the generator problem described in Figure 2.

The projected area of an object rotating about the X or Y axis onto another object in the XY plane is a function of the area of overlap between the two objects (as if both were in the XY plane) and the cosine of the angle between the plane of the rotating object and the XY plane. We independently model the effects of rotation about the X and Y axes, and use qualitative addition to determine the net change in the projected area. We model the extent of rotation about the X and Y axes, respectively, through two angles measured counterclockwise from the XY plane to the plane of the object: θ_Y, the angle between the plane of the object and the Y axis, and θ_X, the angle between the plane of the object and the X axis.

When θ_X or θ_Y is equivalent to 90 or 270 degrees, the projected area of an object onto the XY plane is zero. These angles also mark the point at which the orientation of the object changes. During counterclockwise rotation about the X-axis, when θ_Y is in the interval 90 degrees to 270 degrees, the identities of the topmost and bottommost points will be reversed. During rotation about the Y-axis, the identities of the leftmost and rightmost points can be reversed. In Figure 4, considering object B, if the object has a value for θ_Y in the interval 90 to 270 degrees, then the orientation of edge E1 would have a positive X-component and a negative Y-component since vertex V1 would be a bottommost point instead of a topmost point.

5.4.1 Rotation in the XY Plane

Rotation about the Z-axis can only change the orientation of an object, and not the projected area onto the XY plane. Given, as in Figure 4, that vertex V1 is the topmost point of an object rotating in the counterclockwise direction, and that vertex V2 follows V1, V2 will be the next topmost point. The difficult problem is to determine which extremal point will change first. Given that V2 is the rightmost point and that V3 follows V2, the question now becomes will V2 become the topmost point before V3 becomes the rightmost point? Since the bottommost and leftmost points may also change, there is always a four-way race condition.

For a general polygon, we would require quantitative knowledge (e.g., the internal angles of the polygon) to avoid intractable branching during a qualitative simulation. However, for such commonly encountered special shapes as circles, rectangles and parallelograms, qualitative solutions do exist. For example, rectangles have the property that all four extremal points change simultaneously, eliminating the race condition.

6 Summary and Conclusions

We discussed two themes related to the use of simplifying abstractions to approximate the actual shapes of objects in theories of qualitative spatial reasoning. First, the use of spatial abstractions could hide certain real distinctions. To preserve soundness, abstraction methods should be chosen such that conclusions about spatial relations that are guaranteed to be true can be distinguished from those that the abstractions satisfy, but which the underlying objects need not satisfy.

Second, although many theories have been developed for reasoning qualitatively with shape abstractions, they fail to provide a qualitative model for the abstractions themselves, and although left unspecified, appear to depend on a quantitative description of the abstraction, and thus, the spatial state. This defeats two purposes for qualitative reasoning: to reach useful conclusions in the presense of incomplete knowledge about the world, and to provide an alternate method when reasoning with quantitative knowledge is unnecessarily complex.

We presented a qualitative spatial representation based on the extremal points of an object. This representation allows directional and topological spatial relations to be described through inequality relations on extremal points. Orientation information is inferred directly from knowledge of the connectivity of an object and its extremal points. The representation can be described and maintained in a qualitative fashion, even if objects are allowed to translate and rotate. The representation has been used in a magnetic fields problems solving application [Rajagopalan 94] and in a diagram and text interpretation application [Rajagopalan 94a].

7 Acknowledgements

I thank Professor Benjamin Kuipers for his many suggestions in developing the ideas presented in this paper.

8 References

[Abella and Kender 93] Abella, A. and J. Kender. Qualitatively Describing Objects Using Spatial Prepositions. In: Proceedings Eleventh National Conference on Artificial Intelligence (AAAI-93), 1993.

[Cui, Cohn, and Randell 92] Cui, Z., A. G. Cohn, and D. Randell. Qualitative Simulation Based on a Logical Formalism of Space and Time. In: Proceedings Tenth National Conference on Artificial Intelligence (AAAI-92), 1992.

[Forbus, Nielsen, and Faltings 91] Forbus, K., P. Nielsen, and B. Faltings. Qualitative Spatial Reasoning: the CLOCK Project. *Artificial Intelligence* (1991) **51**, 417-472.

[Freksa 92] Freksa, C., Using Orientation Information for Qualitative Spatial Reasoning. In: A. Frank, I. Campari, and U. Formentini, Eds., *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space*, Springer-Verlag, Berlin, 1992.

[Galton 93] Galton, A., Towards an Integrated Logic of Space, Time, and Motion. In: Proceedings Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93), 1993.

[Herskovits 85] Herskovits, A. Semantics and Pragmatics of Locative Expressions. *Cognitive Science* (1985) **9** 341-378.

[Joskowicz and Sacks 91] Joskowicz, L., and E. Sacks. Computational Kinematics. *Artificial Intelligence* (1991) **51**, 381-416.

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

[Kuipers and Levitt 88] Kuipers B., and T. Levitt. Navigation and Mapping in Large-Scale Space. *AI Magazine* (1988) **9**(2), 25-43.

[Landau and Jackendoff 93] Landau, B. and R. Jackendoff. What and Where in Spatial Language and Spatial Cognition. *Behavioral and Brain Sciences* (1993) **16**, 217-265.

[Maass 94] Maaß, W. From Visual Perception to Multimodal Communication: Incremental Route Descriptions. *AI Review Journal* (1994) **1-3**, forthcoming.

[Mukerjee and Joe 90] Mukerjee, A. and G. Joe. A Qualitative Model for Space. In: Proceedings Eighth National Conference on Artificial Intelligence (AAAI-90), 1990.

[Nielsen 88] Nielsen, P. A Qualitative Approach to Mechanical Constraint. In: Proceedings Seventh National Conference on Artificial Intelligence (AAAI-88), 1988.

[Rajagopalan 94] Rajagopalan, R. A Model for Integrated Spatial and Dynamic Reasoning about Physical Systems. In: Proceedings Twelfth National Conference on Artificial Intelligence (AAAI-94), 1994.

[Rajagopalan 94a] Rajagopalan, R. Integrating Text and Graphical Input to a Knowledge Base. In Working Notes: AAAI-94 Workshop on Integration of Natural Language and Vision Processing, 1994.

[Resnick and Halliday 88] Resnick J. and D. Halliday. *Fundamentals of Physics*, 3rd. Edition, John Wiley and Sons, 1988.

[Retz-Schmidt 88] Retz-Schmidt, G. Various Views on Spatial Prepositions. *AI magazine*, (1988) **9**:2, 95-104.