Efficient Qualitative Kinematics

Boi Faltings, Emmanuel Baechler and Sun Kun
Laboratoire d'Intelligence Artificielle, Département d'Informatique
Ecole Polytechnique Fédérale de Lausanne (EPFL)
MA-Ecublens, 1015 Lausanne, Switzerland
Tel.: +41-21-693-2738, E-mail: faltings@eldi.epfl.ch

December 7, 1990

Abstract

A theory of *qualitative kinematics* is required for qualitative reasoning about the motion of objects in space. Such reasoning is important for many physical problem-solving tasks, such as design or robot planning. The practicality of earlier work has been limited by the complexity and non-compositionality of computing with configuration space constraints. This paper describes a much more efficient and compositional way for qualitative kinematic prediction.

The method is based on a combination of earlier work on kinematic topology ([4]), qualitative mechanism kinematics ([3]) and subsumption detection. We show how each subtheory allows certain kinematic predictions to be made in a largely qualitative manner based on different object models. A complete qualitative kinematic analysis is obtained by combination of results obtained in the different models. This example shows how using several models of the same objects in parallel can yield significantly more powerful and efficient methodologies than were previously possible.

Area: Qualitative Physics **Subarea:** Spatial Reasoning

Submissions: Submitted only to IJCAI-91 and the 5th International Workshop on Qualitative Physics

1 Introduction

There is much interest in spatial reasoning about the behavior of systems of rigid physical objects ([1, 3, 10, 11, 15]). Qualitative reasoning about motion of physical objects is difficult because of two aspects. First, it requires reasoning about kinematic interactions of often complex object shapes, which are difficult to formalize as logical rules. Second, not all interactions are local: moving a piano is difficult because it can be constrained by two distinct objects at the same time. It is promising to base logical spatial reasoning not on elementary shape descriptions, but on an intermediate representation which is computed using numerical methods. In earlier work, we have developed the concept of *place vocabularies* as such an intermediate representation ([3, 6]).

It has been realized by many researchers ([3, 10, 11]) that *configuration space* is a useful formalization of kinematic behavior. In configuration space, geometric constraints are transformed into constraints on the motion of a point. Each possible contact point between objects defines a *constraint* on their motion. Even for simple polygonal objects, the constraints are nonlinear curves in configuration space. The regions of legal motions are traditionally computed by explicitly finding the *envelope* of the set of constraint curves ([12, 14]). This technique is appealing because of its mathematical simplicity, but it turns the original problem of kinematic analysis into an instance of the more general problem of finding a region structure of algebraic curves. Besides the computational complexity, another problem is that the computation is not local: any prediction requires analyzing the complete configuration space, and complete information about all object shapes.

In many domains, the computational complexity and the assumption of complete information make this approach impractical. Consider for the ex-

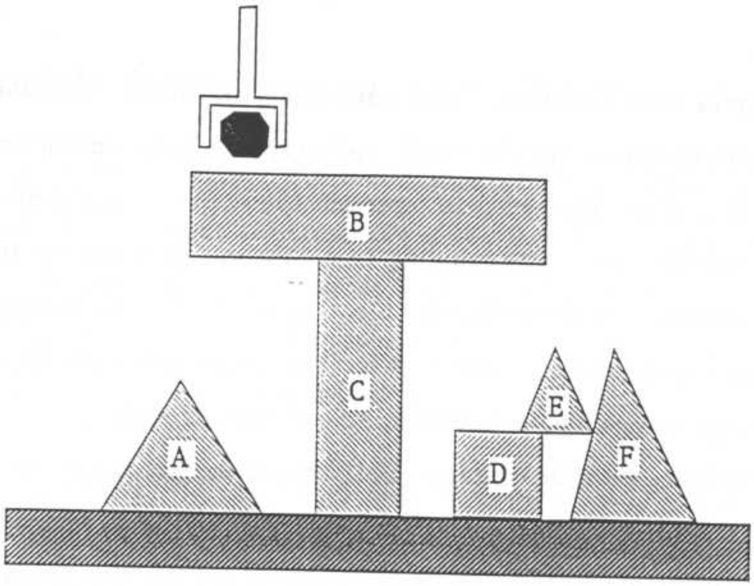

Figure 1: *A prediction problem for a robot: what are the possible consequences of putting down the weight?*

ample the situation shown in Figure 1, where a robot has to determine the possible consequences of putting down an object on top of a pile of others.

In the traditional configuration space formalism, computing the qualitative behavior requires complete and precise information about all object shapes. However, in practice precise information is only required of those objects whose contacts influence the *actual* qualitative behavior. For example, details of the shapes of objects D, E and F in Figure 1 are unimportant for the qualitative prediction of behavior, since the block B will clearly tip to the side of object A. Furthermore, problems with many movable object such as Figure 1 require a configuration space of high dimensionality, in which computation with algebraic curves is intractable.

In this paper, we show a new method for computing the configuration space structure which eliminates these problems. Our algorithm allows incremental computation of only those details of configuration space which are important at some stage in the simulation. Furthermore, it is much more efficient than known methods of computing with configuration spaces for kinematic analysis.

The key idea of our method is to represent the shapes of physical objects by different models, each corresponding to a different view on the physical reality it represents. A region-based object model is used to compute the *topology* of the configuration space ([4]). The qualitative *geometry* of the configuration space region boundaries is obtained using boundary-based object models which define possible contacts. Both sets of predictions are combined with *subsumption* models into a *place vocabulary*, a complete representation of qualitative kinematics ([3]). The process avoids entirely any computation with algebraic curves and the associated computational complexity problems. Furthermore, local predictions can be made using only the information that is actually relevant to the prediction. This allows reasoning with incomplete detail information about object shapes.

We focussed our research on two different problems types: the analysis of higher kinematic pairs with a total of two degrees of freedom, and the motion of an polygon in two-dimensional space with three degrees of freedom. The first example shows that our method is capable of reasoning about complex interactions, while the second shows how to deal with high-dimensional configuration spaces in an efficient manner. Our implementation handles all examples of problems in these classes, and the central aspects of the theory are general beyond these example domains. For example, it will be possible to compute the kinematics of objects with arbitrary shapes by providing the proper extensions to geometric and subsumption models. On the other hand, the extension to three dimensional objects has not been tackled yet.

2 Models for Kinematic Prediction

A *model* of a physical objects is a representation which characterizes its properties. In our system, important properties of the objects and their interaction are captured in three different symbolic models. In this section, we describe each model in terms of the object properties it represents, and the kinematic predictions that follow from these properties.

The Topological Model If the topology of a space is a graph of adjacent regions, the topology of a shape is a graph of adjacent shape primitives. In our representation, we distinguish *pieces* and *cavities* as shape primitives ([4]). A piece is centered on a convex vertex and contains the area of the object behind this vertex and its adjacent boundary edges. A cavity is centered around a concave vertex and contains the empty space ouside of this vertex and its adjacent boundary edges. A cavity also exists around the point where an edge comes closest to the object's center of rotation ([4]). The topology of an arbitrary polygon shape can be represented as a circular sequence of pieces and cavities representing the primitives around the object. The concepts of pieces and cavities can be generalized to sections of convex and concave curvatures, as described in ([4]). A sample decompsition of a polygon is shown on the left in Figure 2.

The kinematic topology of the interaction of a pair of objects is given by the *topology graph*, shown for example on the right in Figure 2. Nodes in the topology graph are either ([4]):

- *obstacles*, representing the region of illegal configurations corresponding to an overlap between two pieces.
- *bubbles*, representing the region where a piece falls in the legal region

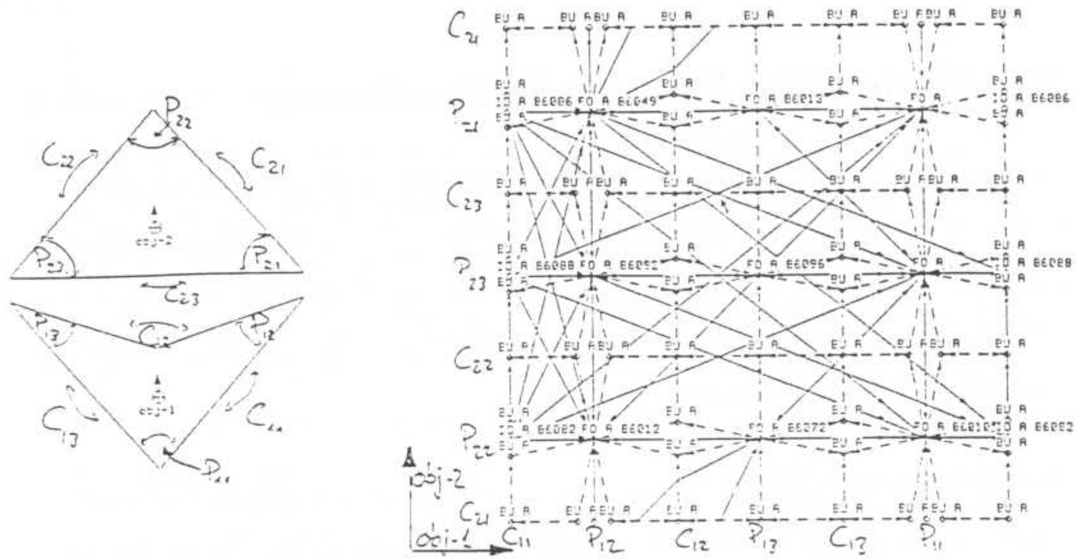

Figure 2: *A topological object model with pieces (P_{ij}) and cavities (C_{ij}). Each combination of primitives creates a bubble (B) or obstacle (FO) in the topology graph.*

represented by a cavity.

- *infinite obstacles*, representing the constraint imposed by the boundary at the bottom of a sequence of cavities.

Adjacencies in the topology graph are given by the adjacencies of primitives on the objects themselves: adjacent shape features generate adjacent contact configurations. In addition, all obstacles have direct connections to those generated by the closest primitives on each of the objects. The *initial* topology graph is laid out on a two-dimensional grid in which each dimension represents a march around one of the objects. Because each object boundary is closed, the graph is circular in both dimensions and forms in fact a closed surface.

A remarkable fact about the initial topology graph is that because it makes no assumptions about directions of motion, it is independent of the actual degrees of freedom of a pair of objects. The *raw* topology graph, which describes the actual topology of a configuration space, is obtained by embedding the intial topology graph in the space as follows.

When motion is unrestricted, there are three degrees of freedom between two objects $\mathcal{O}_1$ and $\mathcal{O}_2$, thus a three-dimensional configuration space. The topology graph then describes the topology of the *surface* of the three-dimensional configuration space obstacle that $\mathcal{O}_1$ poses to the motion of $\mathcal{O}_2$. The raw topology graph of the full three-dimensional configuration space is obtained by linking all nodes of the initial topology graph to a single bubble representing the exterior configuration space.

When motion is restricted, the actual configuration space is only a subspace and parts of the topology graph become inaccessible. For example, when both objects can only rotate, some of the links and obstacles disappear and some of the bubbles are divided into two disjoint pieces. The modifications to the topology graph which reflect this are determined by distance comparisons taking into account the freedom of motion, as described in [4]. Finally, we note that the embedded raw topology graph is often drastically modified by the existance of subsumptions. These modifications are discussed in section 3.

The Geometric Model The geometry of a polygon is given by the edges of its boundary, so the geometric model describes the object as a sequence of boundary edges and vertices with precise coordinates. It is the *metric diagram* ([6]) of our system. Additionally, all parameters relating to the freedom of motion are precisely defined. For example, in the case of two rotating objects, the distance between the centers of rotation is part of the geometric object model.

The geometric model defines the different possible *contacts* which are possible between objects. Two contact configurations are qualitatively equal if they involve the same boundary elements, the qualitative directions of motion ([3]) allowed are the same, and it is possible to pass from one to the other by traversing only qualitatively equal configurations. The geometric model also serves to compute properties of particular object contacts. In contrast to the kinematic topology, these predictions are not independent of the actual freedom of motion. As examples, we distinguish the analysis for unrestricted motion of a polygon among obstacles (3DOF) and that for two objects with rotational freedom only (2DOF).

Three important types of kinematic predictions are made based on the geometric model: possible contacts, local contact adjacencies, and qualitative directions of motion. The set of possible contacts between a pair of polygons is the set of combinations of a boundary elements of one object with a boundary element of the other object. When both boundary elements are vertices, we call such a configuration a *touchpoint*, and use the same term to refer to the curve corresponding to this contact in the 3DOF case. A contact between a vertex and an edge defines an *edge contact*, and a contact between two edges is called a *local subsumption*. Finding the potential contacts is a straightforward combination of the sets of boundary elements.

Every touchpoint and every local subsumption defines an adjacency between four different edge contacts (Figure 3). In any given configuration, only two are adjacent to a touchpoint or a local subsumption. For a touchpoint, the adjacent pairs are computed by comparisons of the angles between edges that meet at the vertices. For a local subsumption, the distinction amounts to a comparison of relative lengths of object edges. In the case of 2DOF, only one pair is possible, which is chosen based on consideration of the object's relative attachments.

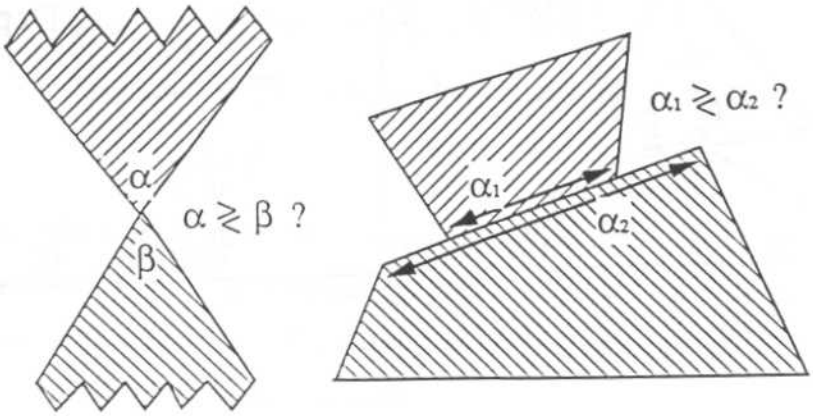

Figure 3: *Touchpoints (left) and local subsumptions (right) define local adjacencies between object contacts.*

The essential feature of kinematics is that a particular contact configuration between objects restricts the possible relative object motions. In order to be useful for predicting behavior, every contact must be broken up according to the relative qualitative motions it allows. This is the third class of predictions in the geometric model.

Qualitative motion is represented using qualitative vectors, as described in [3]. A qualitative motion vector consists of the signs of the configuration space parameters. For two rotating objects, it is a pair of directions of rotation $(\delta\phi, \delta\psi)$. For an polygon moving in space, it is a triple $(\delta x, \delta y, \delta\theta)$ giving the direction of motion of the object center in a Euclidean coordinate system, and the direction of rotation around the center.

In the 2DOF case, a contact fixes a particular qualitative motion vector such that only motion along this vector or its inverse can maintain the contact ([3]). In the 3DOF case, a similar relation holds in touchpoints and subsump-

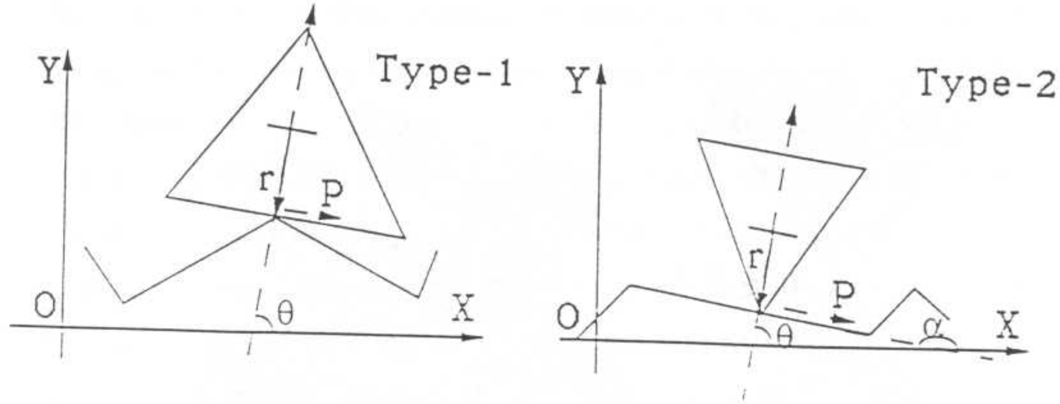

Figure 4: *The two* dead point *contact configurations. Qualitative directions of motion are defined in the text.*

tions, which leave the object with one degree of freedom. For simple contacts, it is possible to state a confluence

$$\pm[\delta x] \pm [\delta y] - [\delta\theta] = 0$$

where $[x]$ refers to the sign of x. Contact configurations with different possible directions of motion are separated by configurations where one of the elements of the consistent motion vector becomes zero, or respectively the confluence changes value. We call such configurations *dead points* ([3]). There are two types of such dead point configurations, shown in Figure 4. We first define a double representation for the orientation angle θ of the moving object as either a qualitative vector $([\theta_x], [\theta_y])$ or as an angle $[\theta]$. The following confluence relates the two notations:

$$[\theta_y][\delta\theta_x] - [\theta_x][\delta\theta_y] - [\delta\theta] = 0$$

Depending on which region the contact point falls in (see Figure 4), the following confluences hold in the contact:

$$\text{Type 1:} \quad [\theta_x][\delta x] + [\theta_y][\delta y] - [P][\delta\theta] = 0$$

$$\text{Type 2:} \quad [\alpha_y][\delta x] + [\alpha_x][\delta y] + [\theta - \alpha + \pi/2][\delta\theta] = 0$$

Note the remarkable fact that the dead point configurations are independent of the choice of (x,y) coordinate system! It is thus possible to use the same place vocabulary with different coordinate systems without extensive recomputation, an important aspect in problem solving.

For subsumptions and touchpoints, all confluences which characterize the possible directions of motion for both contacts hold simultaneously. Since the directions of motion are always unambiguously defined, we can use the qualitative Gauss rule ([2]) to obtain confluences with only two elements. These confluences define the qualitative directions of motion consistent with maintaining both contacts.

Subsumption Model Subsumption models are very different from the topological or geometric models. They are not defined for individual objects, but only for *pairs* of objects. Furthermore, they do not characterize local elements, but the characteristics of a particular subsumption configuration of the two objects, such as shown on the right in Figure 5.

Every pair and triple of possible contacts defines a *potential* subsumption configuration. The actually existing configurations are found by applying several filters to reduce the number of candidates, and the computing the exact configuration by algebraic analysis. Using algebraic transformations ([9]), the system of nonlinear equations characterizing the intersection can be transformed into a single polynomial whose roots correspond to solutions of the

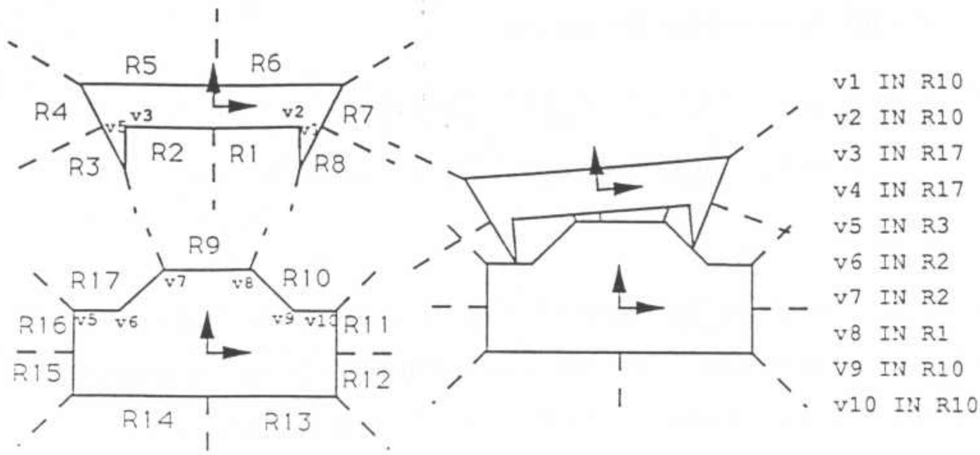

Figure 5: *A subsumption is characterized by several simultaneous unconnected points of contact (right). The subsumption model, which exists whenever such a configuration is possible, gives a table of associations between topological elements which holds in the subsumption configuration.*

system of equations. Depending on the degrees of freedom, the polynomial has 1 or 2 variables, corresponding to a point and a curve in configuration space, respectively. Subsumption models could also be obtained by imagery techniques, a promising solution when shapes are not precisely represented.

Besides establishing an adjacency between particular contact configurations, a subsumption also has profound effects on the kinematic topology of the device. It requires creating a direct adjacency link between two topological objects which may be in very different portions of the graph. Depending on the embedding of the graph, such a link may be inconsistent. We therefore require that any subsumption link must pass through a sequence of elements of the topological graph. These elements are identified in the subsumption model as correspondences between regions defined by cavities on one object, and shape primitives of the other object. This association is shown on the right in Figure 5.

3 Combining Predictions

Local predictions about possible contact relationships are of little use for predicting dynamics unless the adjacencies between them are also known. Furthermore, many potential contacts are in reality impossible because they interfere with other contacts. It is therefore necessary to integrate the local predictions in a global structure, in our case the kinematic topology graph. As a result of this integration, we obtain a full place vocabulary ([3, 6]) representing the kinematics of the device.

Integrating Global Subsumptions with Topology Before the topology graph can be used as a substrate for a place vocabulary, it must be corrected by incorporating the global subsumption predictions. We assume that the legality of the subsumption configuration has already been established at the time the subsumption model was computed.

The subsumption model itself defines the elements of the topology graph which are affected by the link which is to be installed. The modifications to each affected element depend on its type. A bubble of free space connected only to bubbles or only to obstacles is split in two, and a link is created which has the two bubbles on its sides. In the case where the graph is embedded in three dimensions, these two bubbles remain circularly connected around the link. A bubble with free space on one side is left unmodified, and the subsumption link passes through the links between the neighbouring obstacles. These local criteria define completely how a subsumption link is installed in the topology graph.

Integrating Geometric Predictions with Topology Edge and vertex contacts correspond to configurations on the boundary between free and blocked configuration space, modelled topologically as the boundary between obstacles and bubbles. In a consistent solution, every qualitatively different contact corresponds to one piece of this boundary. In particular, connections between obstacles in the topology graph, which we call LINKs, must correspond to connections between contacts, called JOINs. The consistency of the network of LINKs and JOINs together with the correspondences among elements of the geometric and topological object models eliminate all the physically impossible object contacts.

The contacts associated with a particular obstacle are given by the correspondences between topological and geometrical models. In order for a contact to be physically feasible, it must be possible to integrate it consistently with neighbouring contacts. Possible JOINs between contacts are defined in the geometric model by subsumptions, touchpoints and qualitative direction changes (deadpoints). Subsumption models contribute additional JOINs.

The first consistency constraint is that every valid JOIN must be associated with one side of a LINK which it topologically directly adjacent to free space bubbles. The second constraint is that JOINs and contacts must form a consistent surface. The two constraints suffice to determine a unique set of actually valid contacts and adjacencies between them. As each JOIN unambiguously defines two adjacent contacts, it is not even necessary to propagate constraints beyond local adjacency. Contacts adjacent to a particular starting point can thus be computed based only on the complete topology, subsumption models, and information relevant to the starting contact and its neighbours.

No precise information about other parts of the shapes is required.

The Place Vocabulary Once the adjacencies between contacts have been determined, the place graph is constructed by turning all bubbles adjacent to valid contacts into free-space regions and linking them to the adjacent edges.

Finally, each adjacency in the place graph is labelled with the qualitative directions of motion which are consistent with its traversal. We first consider adjacencies which correspond to changes in contact. Such adjacencies are traversed when the point of contact moves beyond the endpoints of a boundary segment. In the case of two degrees of freedom, it turns out that the qualitative direction of motion of the contact point on an edge can be unambiguously inferred from the qualitative direction of motion. In the case of three degrees of freedom, the direction of movement δP of the contact point along a boundary edge (as defined in Figure 4) is characterized by the following confluence:

$$[\alpha_x][\delta x] + [\alpha_y][\delta y] + [\delta\theta] - [\delta P] = 0$$

Thus, we can expect considerably more transition ambiguities in the case of three degrees of freedom. A dead point of type 1 is defined by a particular contact point on the object boundary, so the directions of motion consistent with a transition over it can computed similarly to transitions to other contact configurations. A dead point of type 2 is defined only by the angle between the objects, and so a transition over it is always consistent with any angle change in the proper direction, independently of movements in the other coordinates.

Finally, the directions of motion consistent with breaking or establishing a contact are given by all the qualitative directions of motion that fall within the half plane established the normal vector to the corresponding surface in configuration space ([3]). In the case of three degrees of freedom, the normal

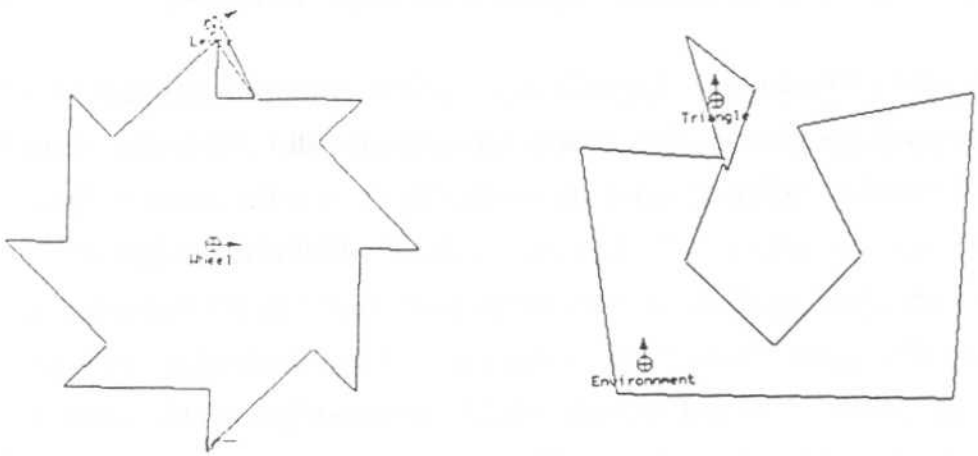

Figure 6: *Two examples that our program has analyzed: a ratchet (left), and a triangle moving in a space of obstacles (right).*

vector is the normal vector to the applicable confluence (as defined in Section 2), whose components are the coefficients of the three quantities.

4 Examples of Problem Solutions

The techniques described in this paper have been implemented for two example domains: a pair of rotating polygons, and a polygon moving freely in a space of obstacles. The two examples, shown in Figure 6, pose different difficulties. In the ratchet example, the difficulty is that the topology graph is embedded on a lower-dimensional surface, as described earlier in the paper. Predictions from the geometric model are also restricted to this lower-dimensional surface. We do not reproduce the results of the complete analysis here, as the resulting place vocabulary is the same as that already shown in [5]. It has been shown to be a sufficient basis for computing an envisionment of qualitative behavior

([13]) in the framework of qualitative process theory ([7]).

The example of the polygon is similar to that analyzed by Davis in [1], or by Forbus in FROB ([8]). The difficulty is that the configuration space is now three-dimensional, making it impossible to rely on planar geometry to compute its characteristics. Place vocabularies for three-dimensional configuration spaces turn out to be extremely large: even for the simple example in Figure 6, the complete place vocabulary has as many as 1300 places! A total envisionment of such an underconstrained system is not likely to be a useful end result. More realistically, one would like to answer questions such as:

- What positions can the object end up in from a given starting position?
- Is a particular combination of contacts possible?
- Are there alternative behaviors to an intended behavior?

The answers to such questions, can be obtained by analyzing a qualitative simulation or the place vocabulary directly. In the case shown in Figure 6, we can for example deduce that all the final positions shown in Figure 7 are possible when the object starts with no motion in the starting position (a), but that no motion will occur when the object starts out at rest in the position (b). While such conclusions are similar to those produced by FROB. they are derived in a more realistic domain which takes into account the kinematic interactions of the moving object with the obstacles. We have thus made a big step towards methods which are able to qualitatively analyze situations such as those shown in Figure 1. We expect the generalization of our methods from three-dimensional configuration spaces to the higher-dimensional spaces involved in such problems to be much easier than the generalization to three dimensions which we have now achieved.

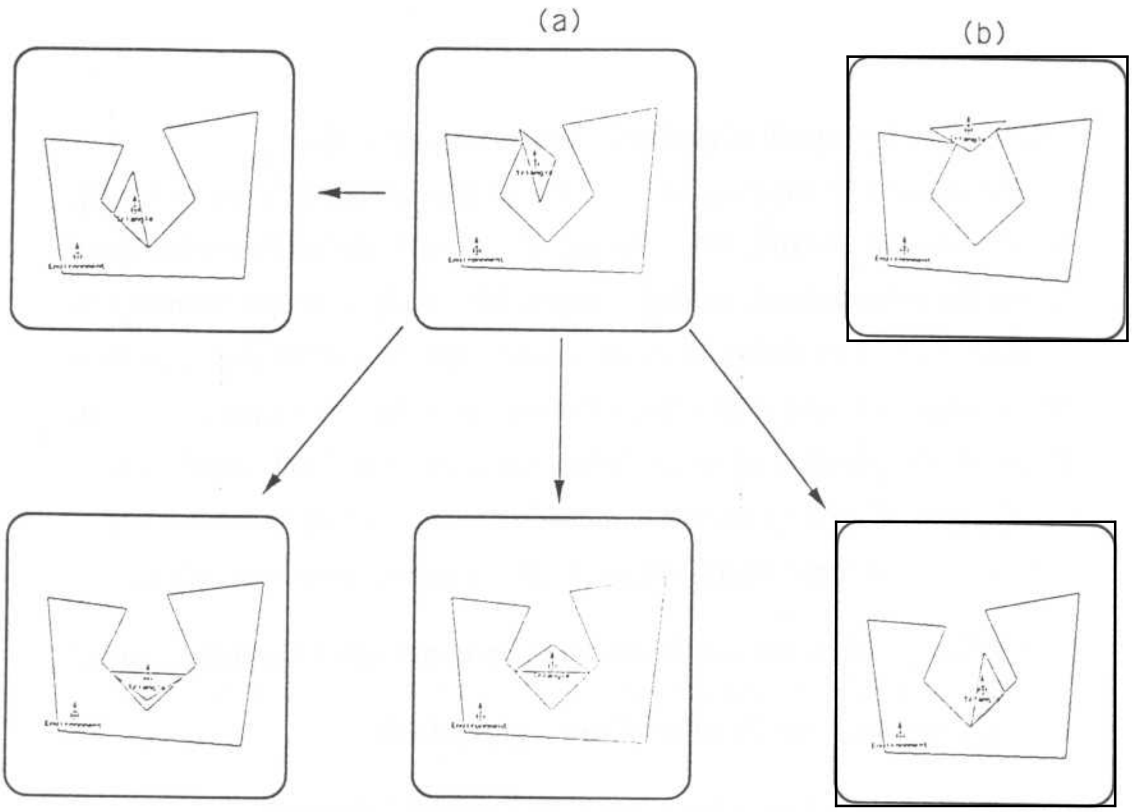

Figure 7: *Different possible behaviors of a polygon under the influence of gravity (downward). Starting at rest from position (a), the object can end up in several final positions as shown. Starting at rest from position (b), no motion will occur.*

5 Conclusions

In this paper, we have seen an example of how the most elegant mathematical formalization of a physical problem can result in unncessary computational complexity of the solution. We developed a less elegant but much more practical solution using different models of the same physical system. The resulting algorithm is not only much more efficient, but can use input information of a very qualitative nature. All numerical computation occurs in the calculation of well-defined models, avoiding the unstructured (and often unstable) computation inherent in classical configuration space analysis.

We believe that similar lessons can be learned in other domains, where a qualitative analysis of the problem can provide the framework within which an intractable mathematical formulation can be efficiently solved. This could result in useful applications of qualitative methods to important practical problems.

We have also made an important step towards making qualitative kinematics practically useful. Our algorithm for computing place vocabularies is far more efficient than methods based on algebraic computations in configuration space. More importantly, we are able to derive characteristics of parts of the configuration space based only on the information relevant to that part. The next steps to be tackled are the extension to arbitrary object shapes and three-dimensional objects. We are also using the algorithms as a basis for automatic mechanical design.

While the techniques we have shown come closer to a purely qualitative kinematics, we are not quite there yet. First of all, the subsumption models are not really compositional qualitative models, but have to be computed anew for each particular pair of objects. Furthermore, several predictions in the geometric model and in the embedding of the topology graph imply numerical computation using dimensions of both objects which can not be expressed as qualitative combinations of properties of the individual objects. Finally, the fact that several models of the same object are used limits the possibilities of backward reasoning from function to shape.

References

[1] **Ernest Davis**: "A Logical Framework for Commonsense Predic-

tions of Solid Object Behavior," *International Journal of AI in Engineering* **3**, 1988

[2] **Jean-Luc Dormoy, Olivier Raiman**: "Assembling a Device," *Proceedings of the AAAI-88*, St.-Paul, August 1988

[3] **Boi Faltings**: "Qualitative Kinematics in Mechanisms," *Artificial Intelligence* **44**(1), June 1990

[4] **Boi Faltings, Emmanuel Baechler, Jeff Primus**: "Reasoning about Kinematic Topology," Proceedings of the 11th IJCAI, Detroit, 1989

[5] **Boi Faltings**: "The Place Vocabulary Theory of Qualitative Kinematics in Mechanisms," Tech. Report UIUCDCS-R-87-1360, University of Illinois, 1987

[6] **Ken Forbus, Paul Nielsen, Boi Faltings**: "Qualitative Kinematics: a Framework," *Proceedings of the IJCAI-87*, Milan, 1987

[7] **Ken Forbus**: "Qualitative Process Theory," *Artificial Intelligence* **24**, 1984

[8] **Ken Forbus**: "Spatial and Qualitative Aspects of Reasoning about Motion," *Proceedings of the AAAI-80*, Stanford, August 1980

[9] **Nathan Jacobson**: "Basic Algebra I," W.H. Freeman and Company, New York, 1985

[10] **Leo Joskowicz**: "Shape and Function in Mechanical Devices," *Proceedings of the AAAI-87*, Seattle, July 1987

[11] **Leo Joskowicz, Sanjay Addanki**: "From Kinematics to Shape: An Approach to Innovative Design," *Proceedings of the AAAI-88*, St.-Paul, August 1988

[12] **T. Lozano-Perez, M. Wesley**: "An Algorithm for Planning Collision-free Paths among Polyhedral Obstacles," Communications of the ACM **22**(10), October 1979

[13] **Paul Nielsen**: "A Qualitative Approach to Rigid Body Mechanics," Tech. Report UIUCDCS-R-1469, University of Illinois at Urbana-Champagin, 1987

[14] **J. Schwartz, M. Sharir**: "The Piano Movers Problem 2: General Techniques for Computing Topological Properties of Real Algebraic Manifolds," *Adv. Appl. Math* **4** (1983), pp. 298-351

[15] **K. Ulrich, W. Seering**: "Function Sharing in Mechanical Design," *Proceedings of the AAAI-1988*, St.-Paul, August 1988