

537-63

7P

118119

N0052545

A Framework for Qualitative Reasoning About Solid Objects

E. Davis
New York University
New York, NY 10012

1. Abstract

Predicting the behavior of a qualitatively described system of solid objects requires a combination of geometrical, temporal, and physical reasoning. Methods based upon formulating and solving differential equations are not adequate for robust prediction, since the behavior of a system over extended time may be much simpler than its behavior over local time. ~~This paper discusses~~ a first-order logic, in which one can state simple physical problems and derive their solution deductively, without recourse to solving the differential equations. This logic is substantially more expressive and powerful than any previous AI representational system in this domain.

2. Introduction

To operate effectively in an uncontrolled environment, an autonomous robot will have to reason about, understand, and predict external physical events. In many circumstances, however, it will be necessary to reason about physical events on the basis of partial information; the objects involved may not be wholly perceived, or the complete physical specifications may be too complex to use, or the robot may need to reason about hypothetical or generic situations. In such cases, the robot will have to reason qualitatively, inferring general characteristics from incomplete knowledge. Human common sense is often very good at speedy prediction of physical events in qualitative terms; conventional computational schemes are typically very poor at it.

Understanding solid objects is particularly important in physical reasoning, and human beings are particularly adept at thinking about solid objects. Our objective is to build an AI program that can reason qualitatively about solid objects and that can derive correct predictions about their behavior in cases where these predictions are intuitively obvious. This is harder than one might first guess, owing to the many complex ways in which the geometry of the objects affects their behavior.

As a first step toward building such a program, we have analyzed the kinds of knowledge needed to support such reasoning, and we have defined a formal language L in which this formal language can be expressed. We have shown that interesting problems can be solved qualitatively by inference from plausible axioms expressed in L . The language L is more expressive and supports richer inferences than any previous representation scheme in this domain. We give the full details of L and its applications in [1]; here, we give only a sketch.

In concentrating on the representation and formulation of knowledge, and postponing questions of algorithms or control structure, we follow Hayes [2]. However, we depart from Hayes' research program in some respects. We do not attempt to model "naive" physics; rather, we have made free use of Newtonian

mechanics, including concepts that have no commonsense analogue, such as total mechanical energy. Also, our proofs are lengthy, violating Hayes' dictum that obvious facts should have short proofs.

The mathematics used here is not "qualitative" in the restricted sense of representing quantities purely in terms of order relationships and constants [3]. Such a representation is too weak to support the inferences needed in this domain.

We have chosen two kinds of problems as foci for our analysis; predicting what happens when a die is dropped inside a funnel (Figure 1) and what happens when a block is dropped onto a table (Figure 2).




Figure 1: A die is released inside a funnel




Figure 2: A block is released onto a table.

Different forms of these problems involve a rich, interconnected body of geometric and physical knowledge for their solution. This paper will focus primarily on the "die in the funnel" example.

3. Background

Several previous AI projects have studied the qualitative physics of solid objects. For example, Fahlman's BUILD program [4] determined the stability of a tower of polyhedral blocks. De Kleer's NEWTON [5] predicted the behavior of a point mass sliding on a

constraints. Bundy's MECHO [6] used force analysis and conservation laws to make physical predictions in situations of specialized format. Forbus' FROB [7] predicted the behavior of a point mass flying among constraints. Funt's WHISPER [8] predicted the behavior of a collection of objects by simulating it in a retina-like image. Novak's ISAAC [9] identified English-language programs of fixed form, and applied special case equations to them. Shoham [10] analyzed the local mobility of an object within constraints.

All these programs provided valuable insights. They were, however, limited in geometrical expressivity and in the range of physics understood. Of these systems, only BUILD dealt with three dimensions; and only MECHO dealt with the motion of extended objects. Only a few kinds of physical interactions were considered.

Another limitation of these programs was more subtle, but more fundamental; they were based almost entirely on extrapolating differential behavior. To make a prediction, the program first determined how each state of the system will tend to change, and then extrapolated these changes to predict a continual trend of change up to the point that the structure of the system changes. This extrapolation could be done qualitatively, as in FROB and NEWTON, or symbolically, as in MECHO, or using point-by-point simulation, as in WHISPER, or by numerical integration, as proposed by McDermott and Berneky (personal communication).

For example, FROB [7] predicts the behavior of a bouncing ball in a well by dividing physical space into 4 regions (the interior of the well, the bottom, and the two sides), and dividing the velocity space of the ball into nine (motionless, up, down, left, right, and the four quadrants.) (Figure 3) There are thus 27 possible states of the system. ($4 \times 9 = 9$ impossible states). The laws of physics are then used to determine which transitions between states are allowed, and thus a transition graph of states is developed. FROB predicts that the system follows a path in this transition graph, ending in a stable state of rest.




Figure 3: Discretized space and velocity in FROB

However extrapolation is done, simulation is inadequate for robust prediction. In this kind of analysis, each different set of boundary conditions is a different system state. Each such state must be separately detected, categorized, and analyzed, and the system's progress through these states must be recorded. Often, however, such a categorization is difficult and pointless. Consider the problem in figure 1; a small die is released inside a large steep funnel. Many states are possible: the die may be in free-fall; it may be colliding or in continuous contact with the top or bottom part of the funnel funnel on any of eight vertices, twelve sides, or six faces; it may be spinning, sliding, or rolling, up, down, or around the funnel. But the prediction that the die comes out the funnel does not require the enumeration of the states and the paths through them.

There are two further arguments. First, the sequence of states traversed depends delicately on the exact shapes, sizes, and physical properties of the die and the funnel, while the conclusion that the die comes out the bottom is robust with respect to small variations in these parameters. Therefore, if the problem is specified with some small degree of imprecision, simulation will either be impossible, or involve some monstrously branching tree of possibilities. But in qualitative reasoning the prediction that the die comes out the bottom should be almost as easy with imprecise data as with precise. Second, the complexity of simulation goes up rapidly with the number of interacting objects. In figure 4, for example, with one die inside another dropped inside a funnel, the set of system states is the cross product of the possible interactions of the two dice with the possible interactions of the outer die and the funnel. Nonetheless, the prediction that the two dice will come out the bottom is intuitively almost as easy as with only one die.




Figure 4: One die inside another released inside a funnel

In short, formulating and solving differential equations is an inadequate technique in this domain, since the behavior of these physical systems over extended time is often easier to characterize than their behavior over local time. A powerful physical reasoning system must be able to infer the general quality of a course of events from broad characterizations of the physical properties of the objects involved, without calculating each subevent.

The programs cited do use some techniques besides simulation. MECHO and NEWTON use energy conservation to prune possible system behaviors. Any state with more mechanical energy than the starting state can be ruled out as a possibility in all future states; for example, the die cannot come out the top of the funnel. FROB predicts that the system ends in a stable state. We believe that effective qualitative reasoning requires more inferences like these, and less use of simulation.

A natural knowledge engineering approach would use rules that state the desired prediction, such as "A small object released inside a steep, large-mouthed funnel will fall out the bottom." But rules of this kind are inadequate, and have rightly been rejected by previous researchers. Any single such rule covers only a small class of problems; covering large classes of problems requires many separate disconnected rules. In particular, a rule like the one suggested above applies only when the die and the funnel are the only objects involved. As soon as another object enters, the rule gives no guidance. That is, such rules are not compositional across objects. Even without other objects, if we allow wide ranges in the shape of the

die and the funnel, the conclusion will sometimes apply and sometimes not. Since there is no simple, general rule for when the die comes out the bottom, a different rule must be stated for each special geometric case.

Maintaining a knowledge base with many special case rules is not effective. First, the knowledge base will have to be large and inefficient. Second, if a new case is not precisely covered in pre-canned categories, the system cannot even begin to deal with it. Third, this approach is aesthetically distasteful. A well-designed system should use similarities among different cases of a die falling through a funnel, and similarities between this problem and similar problems, such as a die shot through a tube, or a die dropped into a box. The analyses of these cases ought to have more in common than the use of rules which are syntactically similar. Finally, it seems plausible that an integrated system of rules will support learning better than a tabulation of special cases.

4. Examples and Analysis

We propose that the "die in the funnel" can be analyzed as follows: (i) Due to the topology of the funnel, if the die goes from inside to outside it, the die must either exit the top or exit the bottom. (ii) Since the die is dropped from rest inside the funnel, it cannot have the energy to exit the top of the funnel. (iii) There is no stable resting point for the die inside the funnel, since it is smaller than the funnel's mouth, and the funnel's sides are steep. (iv) The die cannot stay forever moving within the funnel, for its kinetic energy will eventually be dissipated. Therefore, the die must exit the bottom of the funnel. We claim that in most cases where common sense predicts that the die will come out the bottom, it will be possible to carry out such an analysis, and to support the substeps by inferences from general rules. Different problems will vary in the justifications of the substeps.

Related problems will share parts of the analysis. For instance, in predicting that a die in a small-necked funnel will come to rest at the top of the neck, we may use the identical arguments (i) that the die must either exit the top, exit the bottom, or stay inside; (ii) that it cannot exit the top; and (iv) that it cannot stay inside in a perpetual state of motion. The argument (iii) that it cannot rest stably inside the funnel must be modified to an argument that it can only rest stably at the top of the neck of the funnel; and the additional argument must be made that it cannot exit the bottom of the funnel, since the orifice is too small.

This analysis avoids both problems discussed in section 2. We can avoid analysing, or even determining, the states of motion of the die inside the funnel; all we need to determine is that the die cannot rest stably inside. Different categories of problems are analysed in similar but not identical ways from general principles.

In the rest of this section, we look at variations of this example, and show how this analysis can be applied.

We begin with a simple case (figure 5). The die is a uniform sphere. The funnel is the surface of revolution about a vertical axis of a planar figure with a convex inner side. The radius of the die is less than the radius of revolution of the funnel. The steps of the argument are easily filled in. (i) The top and bottom of the funnel are the only orifices of free space connecting the inside of the funnel with its outside. Therefore, if the die is to go from inside to outside, it must go through the top or the bottom. (ii) Since the die is spherical, its center of mass is in its interior. Since the top of the funnel is horizontal, and directed upward, if the die were to exit it, each point in the interior of the die would be above the top of the funnel at some time. In particular, the center of mass would be above the top at some time. But the die started out from rest below the top of the funnel, and there is no source of additional energy for the die. Therefore, the die cannot come out the top. (iii) By a geometrical argument,

the die can only be inside the funnel in a single point. A uniform sphere can be stably supported at a single point only if the supporting surface is horizontal there. The inner surface of the funnel is nowhere horizontal. Hence there is no resting place for the die inside the funnel.




Figure 5: A spherical die inside a radially symmetric funnel

4.1. Out the top, out the bottom, or stay inside

We now consider how this argument can be generalized and modified. (Further modifications are discussed in [1].) Part (i), that the die must either exit the top, exit the bottom, or stay inside, does not require that the funnel be a solid of revolution; it requires only that the funnel be a tube with only two orifices. We can weaken the condition further, and require only that all orifices other than the top or the bottom be too small to let the die through. Determining whether a die can go through a hole is an easy geometric calculation for various special cases.

4.2. Not out the top


Part (ii), the argument from energy conservation that the die cannot come out the top, depends on the die being convex and on the center of mass of the die starting out below any part of the top. Convexity is only used to establish that the center of mass of the object is in its interior. If this can be done otherwise — for example, by exact calculation, or by establishing that the object shape is a small perturbation of a convex shape — that is sufficient.

A still weaker sufficient condition is that the *ringed filling in* of the die is convex. The *ringed filling in* of a three-dimensional shape S is defined as follows: Consider any planar cross section of S . Let C be any simple closed curve that lies entirely in this cross section. Let p be a point in the plane in the inside of C . Then p is in the *ringed filling in* of S . Figure 6)


Let S be the shape of some object O , and let R be the *ringed filling in* of S . Assume R is convex. Clearly, R is equal to the convex hull of S , so R contains the center of mass of O . Let C be a closed curve lying in S and in a plane containing the center of mass of O . If C goes through a planar surface, then so does every point inside C . Thus, if O exits the top of the funnel, then the center of mass of O must likewise, and the proof goes through. Thus we can establish step (ii) for such shapes as a torus, a wiffle ball, or a cratered convex shape.

4.3. No resting point inside

Part (iii), the argument that the die cannot rest inside the funnel, depended in our first example on the strong




Torus T is cut by plane S



Cross section of T by S
C is a curve in the cross section.
p is a point inside C.

Figure 6

assertions that the die was a uniform sphere and that it could contact the funnel only in a single point. We can easily generalize to nearly uniform, nearly spherical dice. The following formula holds: let θ be the slope of the support; let μ be the coefficient of friction; let ϕ be the maximum angle between the line from the center of mass to a point on the surface and the normal to the surface at that point (Figure 7). The ball can stand still only if $\mu \geq \tan(\theta)$ and $\phi > \theta$. Similarly, if one die is a spherical shell containing another die, they rest stably only if the joint center of mass of the two dice is located directly above the contact point of the outer die with its support, and the inner die rests stably inside the outer die.



N is the normal to the surface.
c is the center of mass.
l is a line through c.
 θ is the angle between N and l.

Figure 7: Distorted sphere

If the die can contact the funnel in several points with different surface normals, the analysis becomes harder. The wider the range of the horizontal component of the surface normals at contact points, the steeper the slope must be, for the normal forces at the various contact points will tend to act against each other, and thus generate larger friction forces. The following rule holds: Let A be in contact with B. Let θ be the minimum slope of the surface of B at a contact point. Consider the horizontal components of the surface normals of B at the

contact points, and assume that there is some direction which lies within some small angle ϕ of all these horizontal components. Let the coefficient of friction be μ . If $\mu < \cos\phi \tan\theta$, then A will slide down B.

Combining all the different ways in which the results (i), (ii), and (iii) may be established, and all the ways in which their geometrical preconditions may be satisfied, gives a rich, interconnected body of results, all with the conclusion, "The die falls out the bottom of the funnel."

5. The Block on the Table

The behavior of the block on the table can be analyzed using a similar argument. After the block is released, it will fall to the table, tipple over a bit, and then move along the table in some combination of sliding, bouncing, and rolling. (Figure 8). It can be estimated how long it will take for the friction involved in sliding and the inelasticity involved in bouncing to consume all the energy gained in the fall and the tipple, and how far the block can travel during that time. A similar estimation can be made for rolling, as long as the object rolls sufficiently poorly. If the surface of the table is uniform, and if these motions will not bring the block off the edge of the table, then it can be predicted that the block will attain a stable state of rest within the estimated time, and within the estimated distance of the point of release.




Figure 8: A block settling on a table

6. The Underlying Knowledge

6.1. Geometry

The arguments in section 3 used several different kinds of geometric knowledge, including:

- The ability to name and describe particular point sets that are connected to objects and are useful for physical reasoning, such as the *top* and *bottom* of a funnel. These are called "pseudo-objects" in our system; the problem of constructing them effectively is the same as the problem of constructing metric diagrams in FROB [8].
- Topological predicates. For example, the funnel forms topologically a box with two orifices, and the die starts out inside the box. [11]
- The use of a property quantified over all irregularities of a certain kind in an object. For instance, the funnel has no holes large enough to let the die through other than the top and bottom.
- Special shapes, such as spheres and surfaces of rotation.
- Inequalities on metric dimensions. For example, the radius of the die is less than the radius of the funnel.

- The bounding of the range of the surface normal over a part of an object's surface. For example, we wish to say that the slope of the funnel is everywhere positive in its inner surface.
 - Convexity and related properties.
- Any adequate geometric language will be strong enough that these, or most of these, can be either expressed directly or inferred.

6.2. Temporal Logic

Our temporal logic follows McDermott's [12]. A *scene* is an instantaneous snapshot of the universe. In our domain, a scene specifies the positions and velocities of all objects. A *chronicle* is a function from the time line to scenes. Chronicles include all continuous motion of objects through space, not just those that are physically possible.

The velocity of an object at an instant is defined to be the limit of its velocity from preceding time. Thus, we can speak of the velocity of an object at the instant of a collision.

The "frame" or "persistence" problem of determining what remains true over time [12, 13] does not arise. There are two classes of predicates in the domain. The first class includes predicates that depend on position and velocity of objects. These are not assumed to remain constant over any interval unless proven to do so. The second class includes structural predicates, depending only on the shapes and material properties of the objects. These are always constant over the problem, and so are defined atemporally. (The closed world assumption is made explicit through the predicate "isolated(OO, C)", which asserts that, during C , no mobile object in the set of objects OO ever comes into contact with any object outside OO .)

6.3. Physics

The world consists of a finite set of solid objects moving in space through time. Objects are rigid and indestructible. The interior of objects may not overlap. Objects have two internal properties besides their shape: a distribution of mass, and a coefficient of elasticity, which determines how the object behaves in a collision. Any pair of objects have a coefficient of friction, which determines the friction forces between the objects.

Objects are subject to four kinds of forces: a uniform downward gravitational force; normal forces, which act to prevent objects from overlapping; friction; and a weak drag force, which dissipates kinetic energy. Certain objects are *fixed*; they do not move, whatever the forces.

Necessary physical deductions include the following:

- Determining whether a set of objects can attain a stable scene while certain geometric conditions hold.
- Finding constraints on the location of the center of mass of an object or a set of objects.
- Resolving a set of forces, and determining motion under those forces.
- Predicting that the existing structure of contacts between objects will change.
- Predicting a collision.
- Predicting the result of a collision.
- Determining whether a chronicle violates a conservation law.
- Characterizing the paths that an object can take without coming to overlap other objects.

7. Ontology

The ontology for our language requires a number of sorts of individuals.

Quantities. Instants of time, quantities of mass, quantities of energy. These are modelled as real numbers.

Points and vectors. These are modelled as elements of R^3 .

Point sets. Subsets of R^3 .

Vector fields. These are functions from some point sets to the space of vectors. For example, the surface normals to an object in a fixed position, directed outward, form a vector field.

Rigid mappings. Mappings from R^3 to R^3 which preserve distance and handedness. These specify a change in position.

General velocities. The derivative of a rigid mapping. A general velocity is the composition of a linear velocity and an angular velocity about a specified axis.

Objects. These are primitive entities. The shape of an object is the point set that it occupies in some particular standard position. This is assumed to be connected, closed, and normal.

Scenes. A scene is a snapshot of the world. Formally, it is a function which maps an object to a pair of a rigid mapping, giving the position of the object, and a general velocity. The place of an object in a scene is image of the object shape under the mapping associated with the object in the scene.

Pseudo-objects. These are point sets that "move around" with objects, like the hole of a doughnut, the opening of a bottle, or the center of mass of any object. Formally, a pseudo-object is a pair of a source object and a point set, designating the point set occupied by the pseudo-object when the object is in standard position. The place of a pseudo-object in a scene is the image of its shape under the mapping associated with its source object in the scene.

Chronicles. A chronicle is a function from an interval of time to scenes.

All chronicles are subject to the following constraints:

- i. All scenes in the range of the chronicle have the same objects in their domain.
- ii. Objects move continuously in space.
- iii. Object velocities are continuous from previous times.
- iv. The velocity of an object is the derivative of its position.

Chronicles do not have to be physically possible. We use the predicate "phys-poss(C)" to distinguish chronicles that obey the laws of physics.

8. Axioms for Physical Reasoning

Based on the above ontology, we have developed a first-order language L and a set of axioms adequate to solve the first "die in the funnel" example. The complete analysis is rather lengthy; the language uses about ninety non-logical terms, not including the standard arithmetic operators, and the analysis involves about 140 axioms. Most (over two thirds) of the terms and axioms are purely geometrical; the rest relate to motion and to physics. We give below three sample axioms, and the complete statement of the "die in the funnel" example, as illustrations.

Geometric Axiom: Smoothness and the value of the surface normal are local properties of the boundary. Specifically, if two bodies share part of their boundary, then, at any interior point of the overlap, one is smooth iff the other is smooth, and their surface normals are either parallel or anti-parallel.

$$\begin{aligned} & [\text{body}(XX1) \cdot \text{body}(XX2) : \\ & \quad XXXA \subseteq \text{boundary}(XX1) \cap \text{boundary}(XX2) : \\ & \quad X \in \text{interior}(XXA) : \text{smooth}(XX1, X)] \rightarrow \\ & [\text{smooth}(XX2, X) : \\ & \quad [\text{surf-norm}(XX1, X) = \text{surf-norm}(XX2, X) : \\ & \quad \quad \text{surf-norm}(XX1, X) = -\text{surf-norm}(XX2, X)]] \end{aligned}$$

Axiom of Motion: If an object O has zero velocity in every scene of a chronicle C , then it stays in the same place throughout C .

$$\begin{aligned} & [\forall S \in \text{scenes}(C) : \text{velocity}(O, S) = 0] \rightarrow \\ & [\forall S1, S2 \in \text{scenes}(C) : S1 \in \text{scenes}(C) : S2 \in \text{scenes}(C) : \\ & \quad \text{mapping}(S1, O) = \text{mapping}(S2, O)] \end{aligned}$$

Physics axiom: The energy of an isolated set of objects OO never increases in a physically possible chronicle C .

$$\begin{aligned} & [\text{phys-poss}(C) : \text{isolated}(OO, C) : T1 < T2] \rightarrow \\ & \quad \text{energy}(OO, \text{scene}(C, T1)) \geq \text{energy}(OO, \text{scene}(C, T2)) \end{aligned}$$

Problem statement: Consider a spherical die, and a radially symmetric funnel. Assume that the inner radius of the funnel is greater than the radius of the die; and that the inner side of a radial cross section of the funnel is convex. If the die is released inside the funnel, and the funnel is held fixed far from the ground, then the die will eventually fall out the bottom of the funnel.

Constants of the example:

odie — the die
ofunnel — the funnel
c — the chronicle
xx-pfunnel — the planar form from which
the funnel is generated
xx-center-line — the axis of the funnel
xcenter — a point on the axis of the funnel

Assumptions:

sphere(shape(odie)).
(The die is a sphere.)

mobile(odie).
(The die is not fixed.)

shape(ofunnel) =
solid-of-revolution(xx-pfunnel, xx-center-line).
(The funnel is the solid of revolution of xx-pfunnel
around xx-center-line.)

planar(xx-pfunnel \cup xx-center-line).
(xx-pfunnel is a radial cross section of the funnel.)

convex-side(inner-dside(xx-pfunnel, xx-center-line),
xx-pfunnel).
(The inner boundary of xx-pfunnel with respect to
xx-center-line is convex.)

distance(xx-pfunnel, xx-center-line) > radius(odie) > 0.
(The radius of the funnel is greater than the radius of the die.)

xx-centerline = make-line(xcenter, vup).
(The axis of the funnel is vertical.)

standard-position(ofunnel, startscene(c)).
(The funnel is oriented in standard position.)

fixed(ofunnel).
(The funnel is fixed.)

isolated({odie, ofunnel, oground}, c).
(The die is isolated from everything but the funnel and the ground. *)

$XF \in \text{shape}(\text{ofunnel}) : XG \in \text{shape}(\text{oground}) \rightarrow$
 $\text{height}(XF) - \text{height}(XG) > \text{diameter}(\text{odie})$
(The funnel is more than the diameter of the die above the ground.)

infinite(c).
(The chronicle is eternal.)

phys-poss(c).
(The chronicle is physically possible.)

motionless(odie, startscene(c)).
(The die starts from rest.)

place(odie, startscene(c)) \subset
tube-inside(shape(ofunnel), s-tube-top(shape(ofunnel)),
s-tube-bot(shape(ofunnel)))

(The die starts from inside the funnel.)

Prove:

exits(odie,
pseudo-object(ofunnel, s-tube-bot(ofunnel, vup)),
c).
(The die exits the bottom of the funnel.)

9. Conclusions

The strengths and limitations of this theory are evident. On the positive side: Using pure first-order logic, we give a formal analysis of a class of problems beyond the scope of any previous AI theory. Our analysis suggests that a qualitative physics for solid objects should include the following features, among others:

- A rich geometrical theory, including topological, metric, and differential descriptors, and special shapes.
- An account of the behavior of physical systems over extended intervals of time. Such an account should incorporate constraints placed by one object on another; conservation laws, especially conservation of energy; the principle that a physical system tends towards a stable resting point; and an account of the net effects of collisions over extended time periods.
- The ability to determine the existence of a stable configuration of objects within qualitatively described geometrical constraints.
- The ability to calculate, exactly or qualitatively, important physical parameters such as the center of mass. [14]
- The ability to bound the effect of small perturbations.

On the negative side: We have not shown that this type of analysis is extensible to cover all, or most, qualitative reasoning in this domain. We have not shown that such an extension would be, in the long run, any more parsimonious than simply enumerating special cases, as in the rule-based method rejected in section 2. We have not shown that any effective computational methods can be developed on the basis of this theory. We cannot give a final resolution to these problems until we have implemented a working system, and determined the range of problems that it is adequate to address.

We plan to begin implementation by developing an adequate geometric representation and inference system. Ultimately, we want to implement a physical reasoning system with all the features mentioned above.

* We need the ground, because otherwise the hypotheses are inconsistent with the axioms. The axioms assert that an infinite chronicle must come to an end in a steady state. Since we show that there is not steady state for the die in the funnel, we must prove it with the ground to rest on.

10. Acknowledgements

Discussions with Sanjaya Addanki, Alex Botta, Pat Hayes, Leo Joskowicz, and Drew McDermott have contributed immensely to the development of this theory. The research has been supported by NSF grants DCR-8402309 and DCR-8603758.

11. References

- [1] E. Davis, "A Logical Framework for Solid Object Physics," NYU Tech. Report #245, 1986
- [2] P. Hayes, "The Naive Physics Manifesto", 1978; revised as "The Second Naive Physics Manifesto", in J. Hobbs and R. Moore, *Formal Theories of the Commonsense World*, ABLEX, 1985
- [3] D. Bobrow, ed. *Qualitative Reasoning about Physical Systems*, MIT Press, 1985
- [4] S. Fahlman, "A Planning System for Robot Construction Tasks," *Artificial Intelligence Journal*, vol. 5, 1974, pp. 1-49
- [5] J. de Kleer, "Qualitative and Quantitative Knowledge in Classical Mechanics," MIT AI Lab Tech. Report #352, 1975
- [6] A. Bundy, "Will it Reach the Top? Prediction in the Mechanics World," *Artificial Intelligence Journal*, Vol. 10, pp. 129-146
- [7] K. Forbus, "A Study of Qualitative and Geometric Knowledge in Reasoning about Motion," MIT AI Lab Tech. Report #615, 1979
- [8] B. Funt, "Problem Solving with Diagrammatic Representations," *Artificial Intelligence Journal*, vol. 13, 1980, p. 201-230.
- [9] G. Novak, "Representations of Knowledge in a Program for Solving Physics Problems," *IJCAI* 5, 1977, pp. 286-291
- [10] Y. Shoham, "Naive Kinematics: One Aspect of Shape," *IJCAI* 9, 1985, pp. 436-442
- [11] E. Davis, "Shape and Function of Solid Objects: Some Examples," NYU Tech. Report #137, 1984
- [12] D.V. McDermott, "A Temporal Logic for Reasoning about Plans and Processes," *Cognitive Science*, 1982, pp. 101-155
- [13] J. McCarthy and P. Hayes, "Some Philosophical Problems from the Standpoint of Artificial Intelligence," in *Machine Intelligence*, B. Meltzer and D. Michies (eds.), Edinburgh University Press, 1969
- [14] A. Bundy and L. Byrd, "Using the Method of Fibres in Mechto to Calculate Radii of Gyration," in D. Gentner and A. Stevens, *Mental Models*, Erlbaum Associates, 1983