

## Chapter 14

# Physical Reasoning

## Ernest Davis

An intelligent creature or automaton that is set in a complex uncontrolled world will be able to act more effectively and flexibly if it understands the physical laws that govern its surroundings and their relation to its own actions and the actions of other agents. In this chapter we discuss work by KR researchers that tries to represent commonsense knowledge and carry out commonsense reasoning over some basic physical domains.

There is, of course, a vast body of computer science and scientific computing which deals in one way or another with physical phenomena; almost all of this lies outside the scope of KR research and hence of this chapter. Even within AI, there are many types of physical reasoning that are excluded here. For instance, the automated visual recognition of a scene is, in a sense, a type of physical reasoning. Image formation is a physical process; the problem in vision is to infer plausible characteristics of a scene given an image of it. Why is this not considered a problem for KR physical reasoning? Mainly because the physics involved is too specialized. A single, quite complex, physical process, and a single type of inference about the process are at issue; and the computational techniques to be applied are highly tuned to that process and that inference and hardly generalize to any other kind of problem.<sup>1</sup>

At the other end of the spectrum, most of the physical theories that appear in the KR literature, such as the STRIPS representations of actions, are too crude and narrow in scope to be of any interest as a physical theory. For instance, the classic blocks world theory applies only to rectangular blocks piled in strict stacks and manipulators constrained to moving a single block from one top of a stack to another; moreover, it does not characterize the positions or motions of the block or manipulator while being moved. The theory is therefore not even a useful start toward a general realistic theory of blocks of general shape in general positions being moved by an actual manipulator.

The most important difference between KR physical reasoning and scientific computing is that, whereas scientific computing almost always aims at achieving a high

---

<sup>1</sup>In principle, high-level physical reasoning could enter into visual recognition, either by providing constraints or measures of likelihood for possible scenes [44] or by relating physical conditions of the image formation process to qualities of the image—e.g., if the lens cap is left on, the image will be black. In practice, the former has been rarely attempted in vision research, and the latter, as far as I know, has never been attempted.

degree of numerical accuracy, KR is almost always content to achieve just a qualitative description. In many cases, predicting qualitative behavior with a high degree of certainty depends on predicting numerical values with a high degree of accuracy—e.g., will the car fall off the cliff, or stop short? In such cases, qualitative reasoning necessarily gives ambiguous results; either the car will stop short and remain intact, or it will fall over the edge and will crash and possibly explode. The quest for numerical accuracy means that most scientific computations involve a fine-grained division of time or space or both (except in the special cases of problems that have an exact symbolic solution). By contrast, KR physical reasoning almost always divides space, time, or space–time into physically significant intervals/regions/histories bounded by significant events/boundaries.

KR also differs from scientific computing in that it often attempts to:

- Incorporate a theory of action.
- Use knowledge for inference in different directions.
- Generate explanations in addition to answers.
- Address everyday domains at the human scale, rather domains that are esoteric, highly specialized, very small or very large.
- Use theories that are psychologically plausible but not necessarily scientifically correct.
- Use explicit theories of causality.
- Study explicitly the interaction between alternative theories at different levels of abstraction. Scientific computation uses many theories at different levels of abstraction, but the problem of choosing the theory appropriate to a situation or of integrating multiple theories in solving a problem is generally left to a human understander (or hard-wired into code).

As contrasted with the *ad hoc* physical theories used in most planning and temporal reasoning, KR work in physical reasoning is distinguished by:

- Generality. The attempt to deal with all or nearly all possible configurations within a given domain. E.g., dealing with arbitrary configurations of blocks of arbitrary shape rather than with stacks of rectangular blocks.
- Continuous time and continuous change over time.
- Geometry and continuous change over space.

Of course, the dividing lines between KR physical reasoning and *ad hoc* KR theories at one end and conventional scientific computing at the other is not a sharp one; indeed, a very important problem for KR is how to integrate all these together.

KR physical reasoning generally involves two important forms of nonmonotonic reasoning. The first is a closed-world assumption, that all the entities that will influence a physical system are known or easily determined. This assumption is made both at the level of theory, that the domain theory accounts for all relevant types of

events, processes, and so on; and at the level of the specific problem, that the problem statement accounts for all the individual objects, actions, and so on. The second is an idealization assumption, that a particular idealization can be safely used. Again this can either be at the level of the choice of theory, such as assuming that the objects in a problem can be modeled as rigid, or at the level of problem description, such as taking a block to be strictly rectangular. Ultimately, it must be expected that KR physical reasoning will have to deal with combining degrees of certainty, and thus require probabilistic or some similar form of reasoning, but little or no such work has yet been done.

Research in KR physical reasoning—which, for the remainder of this chapter we will call simply “physical reasoning”—can largely be divided into four categories:

*Qualitative calculus.* The development of representations and inference techniques for numeric quantities and functions whose value and relations are specified qualitatively. These calculi are the subject of Chapter 9 of this Handbook and are therefore not further discussed here.

*Architecture.* The development of general frameworks which support the statement of physical theories and the description of specific problems and scenarios. Section 14.1 describes the component model and the process model. Again, these theories are presented in Chapter 9, so our description here is brief and focuses on the ontology used in these architectures.

*Domain theories.* The analysis of particular physical domains. Section 14.2 describes kinematic and dynamic theories of solid objects and the theory of liquids.

*Multiple models and levels of abstraction.* Any model of a physical situation used in a reasoning task will include some features of the situation and abstract away others. Thus, a single situation may have many different models, which vary in the features and the detail they include. For instance, depending on the reasoning task, it may be suitable to model a soccer ball as a point object, a perfect sphere, or an irregular sphere; a rigid object or an elastic object; an object of uniform material, a uniform closed rubber shell around an interior of air, or a rubber shell with an inflation hole around an interior of air. Moreover, a reasoner may use more than one of these models in the course of a single reasoning task. The issues of choosing an appropriate model and combining models are therefore critical aspects of physical reasoning. These issues are discussed in Section 14.3.

We conclude in Section 14.4 with a historical and bibliographical survey; here we will mention some further work in the area that falls outside the above categories.

Terminological comment: In this chapter a *fluent* is an entity whose value may change as a function of time. For instance, the fluent “Temperature(O1)” represents the temperature of object O1 as function of time; the fluent “Place(O1)” represents the region occupied by object O1 as a function of time; the Boolean fluent “On(OA, OB)” represents the function of time which is TRUE at times when object OA is on OB and FALSE at other times. A *parameter* is a fluent whose value is in a numeric-valued space, such as temperature. Standard mathematical numerical and geometric functions

are extended to fluents in the obvious way; for instance, if  $f$  and  $g$  are parameters, then  $f + g$  denotes the parameter whose value at any time  $t$  is the sum of the values of  $f$  and  $g$  at  $t$ .

## 14.1 Architectures

An *architecture* for physical reasoning is a representational schema; that is, it is a structure that defines a high-level ontology and a basic set of relations and that supports the representation of various general domains and of specific problems, and the carrying out of particular types of inferences over those representations. Thus, it is roughly analogous to the STRIPS or PDDL representation for planning. The best established and most extensively studied architectures for physical reasoning are the component model and the process model; since these have been already considered in Chapter 9, our treatment of them here is brief and focuses on their ontologies rather than on reasoning methods.

### 14.1.1 Component Analysis

Many complex systems are designed and can be analyzed as a fixed configuration of standard *components*.

A component is an atomic entity with a number of *ports*. Each port has associated with it a number of *parameters* with numerical values. The component imposes constraints on the values of the parameters over time. These constraints are generally either algebraic constraints over the values of the parameters at a given time, or differential equations, relating the derivatives of the parameters at a given time to their values. In the component model, these constraints comprise the entire physical characteristics of the component; aside from the constraints, the component is a black box.

For example, a resistor has two ports  $a$  and  $b$ . Each port  $p$  has two parameters: the inflowing current  $\text{InCurrent}(p)$  and the voltage  $\text{Voltage}(p)$ . A resistor  $r$  is characterized by two equations:

$$\begin{aligned} \text{InCurrent}(a) &= -\text{InCurrent}(b) \quad \text{and} \\ \text{Voltage}(b) - \text{Voltage}(a) &= \text{resistance}(r) \cdot \text{InCurrent}(b). \end{aligned}$$

A capacitor  $c$  has the same types of ports and parameters and is characterized by the equations

$$\begin{aligned} \text{InCurrent}(a) &= -\text{InCurrent}(b) \quad \text{and} \\ \text{InCurrent}(b) &= \text{Capacitance}(c) \cdot \text{Derivative}(\text{Voltage}(b) - \text{Voltage}(a)). \end{aligned}$$

A *node* is a collection of ports connected together. The node imposes a constraint on the parameters of the ports determined by the domain theory. For instance, in the electronics domain, if ports  $p_1 \dots p_k$  are connected at a node, then that creates the constraints

$$\begin{aligned} \text{InCurrent}(p_1) + \dots + \text{InCurrent}(p_k) &= 0 \quad \text{and} \\ \text{Voltage}(p_1) = \text{Voltage}(p_2) = \dots = \text{Voltage}(p_k). \end{aligned}$$

A *system* is defined by a collection of components, and a partition of their ports into nodes. The structure of connections and the component characteristics are fixed over time; what varies over time are the values of the parameters. The set of constraints generated by the components and by the nodes determines the behavior of the system over time.

Electronic systems are the archetypal and best example of a domain that can be analyzed using the component model. The model has also been applied to hydraulic devices, heat transfer systems, and mechanical systems of certain types.

Actions can be incorporated into the component architecture by modeling an agent as an exogenous signal. That is, an agent is modeled as a component for which the values of the parameters are not determined by the theory and the remainder of the system, but rather can be “chosen”. For example, in the electronics domain, an agent could be a voltage source that can choose a waveform to output; the waveform it chooses is its action.

Typical reasoning tasks carried out over component models include:

- *Static evaluation.* If all the constraints are algebraic, then determine the state (or the set of possible states) of the system.
- *Initial value problem.* If the constraints include differential equations, then determine the progress of the system following some starting condition.
- *Signal response.* Determine the output of a system in response to a specified signal at some input.
- *Comparative static evaluation.* Determine the effect of changing some component characteristic on the static state of the system.
- *Comparative dynamic evaluation.* Determine the effect of changing some component characteristic on the dynamic progress of the system.

The best known program using the component model was the ENVISION program of de Kleer and Brown [17]. ENVISION used the sign calculus to solve qualitatively the initial value problem and the comparative static evaluation problem. ENVISION also proposed a model of causality, in which an change to some exogenous parameter in the system causes changes to other parameters by propagating through the network, in a manner that has a sequence, though no measurable time duration.

### 14.1.2 Process Model

In the process model [23], change is brought about by *processes*, *events*, *actions*, and *indirect influences* between parameters.

A *process* is active over a time interval. It is characterized by preconditions and effects. The preconditions must hold for the process to begin, and they must continue to hold throughout the interval in which the process is active. If the preconditions cease to hold, then the process stops. The effects of a process are *direct influences* on numeric fluents. A direct influence is a contribution to the derivative of the fluent; the derivative of the fluent is the sum of the influences of all the processes that act on it.

For example, consider the process of a tap  $t$  filling a bucket  $b$ . The preconditions are that the tap is open, the bucket is under the tap, and the bucket is not yet full.

The process directly influences the fluent “volume of water in the bucket”; that is, the derivative of the volume of water is a sum of terms, one of which is the flow-rate of the tap  $t$ . For example, if there are several taps filling  $b$  and also a leak from the bottom of  $b$ , then the derivative of the volume of water in  $b$  is the sum of the flow-rates of the taps minus the flow-rate of the leak.

An *action* takes place at an instant. It is characterized by preconditions, which must hold for the action to be feasible, and effects, which are discontinuous changes in the value of a discrete or numeric fluent. For example, turning on a tap is an action. The precondition is that the agent is next to the tap and that the tap is closed. The effect is that the tap is open. If the preconditions of an action are satisfied, then an agent has the choice of whether or not he wishes to perform the action.

An *event* is similar to an action except that it is not a matter of choice; it is a natural discontinuous change that must take place if the conditions are met. For instance, suppose that you have a weak bucket whose bottom will fall out when the bucket is half full. Then the event “Bottom of  $B$  falls out” has the precondition that the bucket is at least half full and has the effect that what was formerly a bucket is now a disconnected cylinder and a pan.

Finally, parameter  $p$  is an *indirect influence* on parameter  $q$  if there is a natural constraint relating their two values. For example, the volume of liquid in a bucket is an indirect influence on the height of liquid in the bucket. It is assumed that the system of influences on system parameters can be structured in such a way that (a) no parameter is both directly and indirectly influenced; (b) the relation “ $p$  indirectly influences  $q$ ” is acyclic.

The QP program [23] uses a process model to carry out qualitative projection. Conditions are conjunctions of discrete values, such as “The tap is open” and inequalities, either between one parameter and another, or between a parameter and a constant “landmark” value, such as “The level of water in the bucket is less than the depth of the bucket.” Influences are specified in terms of their sign; e.g., the process of a tap filling a bucket has a positive influence on the volume of water in the bucket, while the process of leaking has a negative influence. Using this information QP can generate an “envisionment graph”, a transition graph between states of the system. Any possible behavior of the system corresponds to a path through the envisionment graph. (The converse does not in general hold; there are often paths through the envisionment graph that do not correspond to physically possible behaviors.)

Both the component model and qualitative process theory are discussed at much greater length in Chapter 9.

## 14.2 Domain Theories

The person on the street is familiar with hundreds, perhaps thousands, of physical categories, qualities, and phenomena; an expert (scientist or engineer) knows perhaps tens or hundreds of thousands; collective scientific knowledge must include many millions. It seems likely that the largest part of achieving general purpose physical reasoning, at either the commonsense or the expert level, will be the representation of all the different concepts involved. To date very few physical domains—certainly fewer than a dozen—has been studied in any depth in the KR literature. In this section, we will look at theories of rigid solid objects and theories of liquids.

### 14.2.1 Rigid Object Kinematics

Solid objects enter into almost all scenarios that physical reasoning in a terrestrial, human-scale environment deals with. More specifically, in a significant fraction of physical reasoning, only solid objects are significant, only the motions of the objects are significant, and the objects can be idealized as rigid (constant shape).<sup>2</sup>

The complete theory of rigid object dynamics is discussed in Section 14.2.2. First, however, we will discuss the *kinematic* theory of rigid solid objects. The kinematic theory is much less informative than the dynamic theory but is nonetheless sufficient in many important applications, and in fact has been applied much more extensively and successfully.

The kinematic theory asserts four rules governing the shape and motion of solid objects:

- The shape of an object is a closed, regular, connected region.<sup>3</sup>
- The shape of an object is constant over time.
- The position of an object is a continuous function of time.
- At any given time, the regions occupied by two distinct objects do not overlap.

In the kinematic theory, therefore, the only significant time-invariant characteristic of an object is its shape, and its only significant time-varying characteristic is its position. The shape can be characterized in terms of the spatial region occupied by the object in some standard position. The position of object  $o$  at time  $t$  can be characterized in terms of a rigid (orthonormal) mapping, characterizing its displacement from its standard position to its position at  $t$  (Fig. 14.1).<sup>4</sup> Thus the kinematic theory can be formulated in first-order logic using the functions  $\text{Shape}(o)$  which maps an object  $o$  to the region which is its shape;  $\text{Position}(o)$  which maps object  $o$  to the fluent of its position over time;  $\text{Place}(o)$  which maps object  $o$  to the fluent of the region it occupies over time; combined with suitable temporal and geometric primitives.

Given a set of objects  $o_1 \dots o_k$  and given the shapes of these objects, a *configuration* is a specification of the position of each object. A configuration is *feasible* if no two objects overlap. A configuration  $c_2$  is *attainable* from configuration  $c_1$  if it is possible to move the objects from  $c_1$  to  $c_2$  without causing two objects to overlap. Given a set of objects and an initial configuration  $c_1$  the *attainable configuration space* is the set of feasible configurations attainable from  $c_1$ . Since the position of objects is a

---

<sup>2</sup>One reflection of the cognitive salience of this category is the persistent attempt in eighteenth- and nineteenth-century physics to reduce all physics to mechanical interactions of small solid objects; e.g., the kinetic theory of heat, or Maxwell's mechanical model of electrodynamics.

<sup>3</sup>A *closed* region is one that includes its boundary. The decision to use a closed rather than an open region is arbitrary, but it simplifies description to specify one or the other. A closed region is *regular* if it is equal to the closure of its interior, and thus is "thick" everywhere and does not have any one- or two-dimensional pieces.

<sup>4</sup>A displacement is a composition of a rotation around the origin and a translation. A translation in  $k$  dimensions is characterized by a vector  $\vec{t}$ ; any point  $x$  is mapped into  $x + \vec{t}$ . A rotation in two dimensions (relative to a fixed origin) is characterized by an angle  $\phi$ . A rotation in three dimensions is characterized by three angles; there are a number of different systems of angles that can be used for this purpose, such as the Euler angles. Alternatively, a  $k$ -dimensional rotation can be characterized by a  $k \times k$  orthonormal matrix.




Figure 14.1: Shape, place, and relative position of a rigid object.

continuous function of time, a configuration  $c_2$  is attainable from  $c_1$  just if there is a path from  $c_1$  to  $c_2$  through the space of feasible configurations for the objects; thus, an attainable configuration space is a path-connected component of the space of feasible configurations. For initial-value problems, in which the shapes of the objects and the initial configuration are given, it suffices to consider only attainable configurations, since no other configurations can ever occur.

Indeed, initial value problems with complete shape specifications can be addressed as follows: One begins by computing the attainable configuration space for the system; that is, the connected component of the configuration space containing the initial configuration. Having done that, the entire content of the kinematic theory lies in the statement that the configuration moves continuously through that space. This technique is particularly effective if the configuration space is of low dimension; that is, the physical system has few degrees of freedom. Significantly, this is often the case with man-made mechanisms; indeed, for many mechanisms, such as gear trains, the configuration space is one-dimensional, or nearly so.<sup>5</sup> In such cases, it is easy to determine the consequences of the constraint that the configuration changes continuously. For example, if the configuration space is partitioned into regions, then the continuity constraint means that the configuration must move between adjacent regions in the space.

A number of methods for qualitative analysis for kinematic systems have been developed. The most common method [20, 48, 50] starts with exact shape descriptions, computes the configuration space exactly, divides the configuration space into significant regions, and then characterizes qualitative properties of the system from the connectivity of these regions. Kim [38] describes a system for qualitative reasoning about linkages, analyzing the relation between the directions between the ends

<sup>5</sup>Man-made mechanisms tend to rely on kinematic constraints when possible, because they are extremely robust. A large external force or impact is generally required to make solid objects significantly bend or break, and there is no way to cause two solid objects to spatially overlap.



of the arms (discretized into quadrants), the angles between the arms (likewise), and inequalities between the lengths of the arms.

A theory of action can be integrated into a kinematic theory by specifying that specified objects are *manipulable*, and that their motions are thus chosen by the agent. In this setting, a standard projection problem consists of a specification of the shapes and initial positions of all the objects and the motions of the manipulable objects. The kinematic theory asserts that the other objects will move through the configuration space along a path that accommodates the specified motions of the manipulable objects, if there is such a path; if there is not, then the specified motions are infeasible. The most difficult aspect of formulating this theory is asserting that an action is feasible unless it leads to an infeasible configuration.

In some cases, it is convenient to abstract a kinematic system using a simplified shape description together with a set of imposed constraints. For example, mechanical systems often contain parts such as gears that are pinned by a circular pin to a fixed frame so that they can rotate around the pin. It is common to abstract away both the frame and the pin, and to view the gear as subject to an abstract constraint that enforces the condition that the center of the gear remains fixed (Fig. 14.2) (e.g., Faltings [20] and Joskowicz [34] use this device for gears rotating on a frame, and Kim [38] uses the analogous device for linkages).

### 14.2.2 Rigid Object Dynamics


The kinematic theory of solid objects, though often very useful, is in general much too weakly constraining for commonsense reasoning. The *dynamic* theory of rigid solid objects describes the motions of solid objects in all circumstances in which they do not break or significantly bend. Thus, for example, the fact that a book remains on a bookshelf rather than floating off into the air, or that a chair will be stable when standing on four legs but not when standing on one leg lie beyond the scope of the kinematic theory; they require at least part of the dynamic theory.

It has been known since the early eighteenth century that the interaction of rigid solid objects is characterized by the following rules: the kinematic principles listed above; Newton's second and third laws; the existence of a normal force between objects at a contact point; static and sliding Coulomb friction between objects at a contact point; and a theory of instantaneous momentum transfer when objects collide. For terrestrial problems at the human scale, these must be supplemented by the existence of a uniform downward gravitational force; the existence of fixed objects (such as the ground) which never move; the existence of manipulators which can be subjected to an applied force at the will of an agent; and a closed world assumption that the only types of forces that act on objects are those enumerated in this theory.


Somewhat surprisingly, there is still no complete, accepted formulation of this theory in the scientific literature, particularly the theory of collisions. Even in the simple case of two objects colliding at a point, there is debate over the proper theory,<sup>6</sup> and

---

<sup>6</sup>The desiderata for such a theory are that it corresponds to experiment; that it satisfies global constraints, such as conservation of energy, momentum, and angular momentum; that it yields a solution for all well-posed initial-value problems; that numerical calculations converge; and that it can be justified in terms of a more detailed elastic model of solid objects.



Concrete: Gears pinned to a frame.



Abstraction: 2D Gears constrained to rotate around a fixed center.

Figure 14.2: Gears and their abstraction.

there is no standard theory to use in either the case of two objects that collide along a surface or a curve, or the case of collisions involving multiple objects simultaneously. Stewart [56] reviews the state of the art in the current theory.

In any case, the scientific theory outlined above is not well-suited to the needs of reasoning in ordinary applications. It involves determining entities, such as forces, which are only occasionally of interest in commonsense reasoning, and it characterizes behavior over differential time, whereas the reasoner is generally concerned with behavior over extended time. For example, if you put a book on a shelf, you are not usually concerned with the forces between the book, the shelf, and the other books; you are only concerned to predict that the book will stay on the shelf. Similarly, if you carry a loose collection of objects in a closed box from one place to another, you are not usually concerned with the forces between the objects during the journey, or even with how the objects shift their relative positions inside the box. Generally, it suffices to determine that the objects remain inside the box throughout the journey.

Though a few AI programs have addressed the general problem of solid object dynamics by doing full numerical simulation (e.g., [29]) most AI program have dealt with restricted special cases:

- *Point objects.* The NEWTON program [16] performed qualitative prediction of the behavior of a point object on a track. The shape of the track was characterized in terms of the signs of its slope and its curvature. This was the first application of the sign calculus in AI physical reasoning. The FROB program [22] similarly performed qualitative predication of the behaviors of a collection of point objects moving in a world with fixed barriers, and one vertical and one horizontal dimension.
- *Statics.* An important category of physical prediction is to predict that an object will remain unchanged: a book will remain on a shelf, a building or bridge will continue to stand. (Note the contrast here with the usual attitude in KR that this can simply be assumed by default.) The equations of motion and their analysis are of course very much simplified if all that is required is to distinguish between situations that have a static solution and those that do not. Fahlman [21] implemented a static analysis of configurations in the blocks world.
- *Quasi-statics.* In a quasi-static problem, objects all move so slowly that their momentum is negligible as compared to the frictive forces acting on them. Hence objects only move while being pushed, directly or indirectly, by an exogenous force such as a manipulator. The standard scenario for quasi-static problems is a collection of flat objects on a horizontal surface being pushed around, though other scenarios are possible (e.g., a collection of three-dimensional objects in a highly viscous liquid). Exact quasi-static predictions were carried out by Mason [43] to carry out “sensorless manipulation”; i.e., finding ways to maneuver objects to a desired target position without any sensory feedback describing the positions of the objects. Qualitative quasi-static predictions were carried out by Forbus, Nielsen, and Faltings [24], and Stahovich, R. Davis, and Shrobe [55] using qualitative representation of configuration space and of the driving forces. If the motions of the objects are highly constrained, then the quasi-static theory is often equivalent to just the kinematic theory plus the default assumption that objects only move when necessary.

As mentioned above, a theory of action can be integrated into a dynamic theory of rigid objects by designating particular objects as manipulators which are subject to exogenous forces chosen by the agent. Thus, one conceptualizes the robot’s hand as a rigid object which, at the robot’s command, fires invisible rockets to exert specified forces on it. The advantage of this model is that it gives a well-formed boundary problem; a problem consisting of a specification of the initial state plus the forces on the manipulators always has a solution [56]. The disadvantage is that this is not usually a very natural way to think about a manipulator. The natural way to think about a manipulator, indeed, depends on the circumstance: often, it is just a geometric specification of the motion of the manipulator, but sometimes it is a force exerted by a stationary manipulator against an object, sometimes, it is the combination of a motion of the manipulator together with a force exerted on an external object, and sometimes, as in compliant motion, it is a control strategy where the force and motion of the manipulator depends on feedback. No general high-level language suitable for commonsense reasoning has been found for this.




Figure 14.3: Nail in a board.

Another difficulty in the theory of the dynamic theory of solid objects is that the theory is sporadically underdetermined. In most cases, a specification of the initial positions and velocities of all the objects and their material characteristics is sufficient to determine their behavior, but there are exceptions, and these exceptions can be difficult to deal with. The most important category of exceptions is configuration in which an object is jammed. For instance, consider a nail in a hole in a board, pointing up (Fig. 14.3). Will the nail fall out of the hole? It depends on whether the nail was placed in the hole or whether it was driven into the hole. In the latter case, there are large, normal forces on the nail from the board and a corresponding large frictional force holding the nail in place. Thus, the boundary conditions in this problem include a specification of the forces, whereas in most cases forces generally determined by the positions and velocities. This makes it difficult to state what constitutes an adequate representation of a situation.

In some cases, considerations of mechanical energy give powerful constraints. For instance, de Kleer's *NEWTON* [16] uses an energy-based calculation to predict whether a roller-coaster on a track will go around a loop-the-loop, slide back, or fall off. Davis [9] shows how energy considerations can be used to construct an argument that a marble dropped inside a funnel will come out the bottom. (It cannot come out the top, because of conservation of energy; it cannot attain a stable resting position inside, because of the slope of the sides; it cannot remain inside forever moving around, because the kinetic energy will dissipate. Hence, the only possibility is that it will come out the bottom.)

KR work to date has barely scratched the surface of a commonsense understanding of this domain. Most commonsense inferences involving solid objects cannot even be represented in current KR theories, much less implemented.

### 14.2.3 Liquids

Liquids are in one way simpler than solid objects; they do not have a fixed shape that has to be represented and reasoned about. Thus, for example, it is often easier to determine whether a liquid will flow out of a tilted cup than whether an object will fall out of a tilted box. If you are tilting a cup of liquid, then the liquid will start to flow over the side of the cup just when, if there were no such flow, the volume of the inside of the cup below the opening would be less than the volume of the liquid. No such simple rule can be stated for tipping solid objects out of boxes.

On the whole, however, liquids are much more difficult to reason about than solids, because they are not individuated into objects. Rather, a system with liquids can be characterized in three complementary ways [33]. The first method is to define fluents  $\text{Volume}(l, r)$ , the volume of liquid  $l$  in region  $r$ , and  $\text{Flow}(l, b)$ , the flow out liquid  $l$  through directed surface  $b$ . (The regions involved need not be fixed regions in space; they can be fluents whose value at an instant is a region, such as “the inside of a pail”, which moves if the pail moves.)

The second method is to define a fluent  $\text{Place}(c)$  which denotes the region occupied by a “chunk”  $c$  of liquid. Note that  $\text{Place}(c)$  may be a disconnected region. A variant on the second method is to fix a starting reference time  $T_0$ , to identify the region  $\text{place}(L, T_0)$  occupied by liquid  $L$  time  $T_0$ , and then to characterize the evolution of the liquid over time in terms of a fluent  $\text{LiquidTrajectory}(X, L)$ . For any point  $X \in \text{place}(L, T_0)$ , liquid  $L$ , and time  $T$ , the value of  $\text{LiquidTrajectory}(X, L)$  at  $T$  is the location at  $T$  of the particle of  $L$  that was at  $X$  at  $T_0$ .

A third approach is to treat the liquid as a collection of molecules or small particles [7, 30, 53, 18], whose positions and velocities can be tracked (if there are few enough) or characterized. The chief difficulty here is to characterize the interaction between particles in such a way as to give the characteristic liquid behavior.


If we exclude from consideration both mixtures of liquids and phase changes such as evaporation, and we assume that all liquids are incompressible, then we can state the following three kinematic properties:

1. A liquid moves continuously.
2. A liquid does not overlap with a solid, nor do two liquids overlap.
3. A quantity of liquid maintains a constant volume.

In a region-based representation, constraints (1) and (3) above are achieved by asserting the divergence theorem that  $\text{Derivative}(\text{Volume}(l, r)) = -\text{Flow}(l, \text{Boundary}(r))$  and that the flow out through boundary  $b$  is the negative of the flow through  $b$  with the reversed orientation. In a chunk-based representation, these constraints are achieved by asserting that  $\text{Place}(c)$  is a continuous function of time for every chunk  $c$  and that  $\text{Volume}(\text{Place}(c))$  is constant over time.

However, unlike the solid case, the kinematic theory of liquids is not by itself strong enough to analyze many interesting physical situations; a stronger dynamic theory must be used. The dynamic theory of liquids is much less well understood than the dynamic theory of solid objects, both in scientific and in commonsense theories. A few special cases are worth noting:

*Statics, bulk liquid:* If we ignore the phenomenon of a liquid wetting a solid surface, then we may state the following rule: If a body of liquid occupies a connected region  $R$  and is at rest, then the boundary of  $R$  must meet solid objects everywhere except at a collection of horizontal upper surfaces of the liquid. If at all such surfaces the liquid meets the open air, then all these surfaces are at the same height. Otherwise, if some of the surfaces meet bodies of gas that are themselves enclosed by solids, then the difference in heights among the surfaces is proportional to the difference in pressure in the bodies of gas involved (Fig. 14.4). (Note that in such cases, it is necessary to represent the gas explicitly, whereas this is not necessary if all bodies of gas are



The heights in A and B are equal because both meet the open air.  
 The pressure of the gas at the top of C is greater than atmospheric pressure,  
 by an amount proportional to the difference in heights.

Figure 14.4: Liquid statics.

connected to the outside air.) In particular, if a volume  $V$  of liquid is poured into a open solid container, then it will reach a height  $h$  such that the volume of the interior of the container below  $h$  is equal to  $V$ .

*Quasi-statics:* If the solid objects that are in contact with the liquid, and with the contained gases that meet the liquid, are all moving slowly, then it is sometimes possible for the liquid to flow in such a way that the above static constraints are maintained. When this is possible, it generally happens. (It becomes impossible when the liquid is poured out from its container.) In such a case, the above static rules can be used to predict the trajectory of the regions occupied by the liquids and gas, and the flow of the liquid, given the motion of the solids.

Kim [38] describes a system that carries out qualitative predictions of the motions of liquids in response to the motions of pistons. She also includes in her model a special case of solids being acted on by liquids, namely the opening and closing of one-way valves.

Hayes [33] identified 15 disjoint and exhaustive physical states of liquids (Table 14.1). Any quantity of liquid at any time can be divided into parts, each of which is in one of these states. Any quantity of liquid, considered over an interval of time, can be divided into *histories*—that is, regions of space–time—each of which is in a single state. Hayes proposed that a qualitative physics of liquids could be developed in terms of axioms describing how different types of histories meet one another and meet histories of solid object trajectories, on both spatial and temporal faces; and he began work on such an axiomatization. For example, the bottom face of a “falling” history must have a downward flow through it. All but the top, horizontal face of a “pool” history must meet the outer face of solid objects. This axiomatic work was never completed (or even extended beyond Hayes’ original article) for a number of reasons, chiefly because a useful theory would require a much stronger spatial language than Hayes originally envisioned.

Table 14.1. The possible states of liquid (from [33])

	Lazy still	Lazy moving	Energetic moving
Bulk on surface	Wet surface	Flowing down a surface, e.g., a sloping roof	Waves lapping a shore (?) jet hitting a surface (?)
Bulk in space	Contained in container	Flowing along a channel, e.g., river	Pumped along pipeline
Bulk unsupported		Falling column of liquid, e.g., pouring from a jug	Waterspout, fountain, jet from hose
Divided on surface	Dew, drops on a surface		
Divided in space	Mist filling a valley	Mist rolling down a valley	Steam or mist blown along a tube
Divided unsupported	Mist, cloud	Rain, shower	Spray, splash driving rain

### 14.3 Abstraction and Multiple Models

A characteristic of physical reasoning, at both the commonsense and the expert level, is the existence of many different theories for a given domain, and many different ways and levels of detail for describing a given situation, and many different abstraction techniques for simplifying problem statements and problem solving. A reasoner faced with a real-world problem must almost always choose among these in formulating his problem; frequently, he must apply different, mutually inconsistent, theories to different parts of the problem-solving process. Some interesting, but very preliminary, studies have been made of the ways in which appropriate theories/descriptions can be chosen and integrated in problem solving.

Some of the more important categories of abstraction include:

*Alternative physical theories.* Two physical theories  $\mathcal{U}$  and  $\mathcal{V}$  of the same physical domain may be related in that

- $\mathcal{U}$  adds additional constraints to  $\mathcal{V}$ ; that is,  $\mathcal{V}$  is logically a subtheory of  $\mathcal{U}$ . E.g., the relation between the dynamic and the kinematic theory of solid objects.  $\mathcal{U}$  is called a “theorem increasing” [31] or “model decreasing” [47] extension of  $\mathcal{V}$ .
- $\mathcal{U}$  adds additional entities to  $\mathcal{V}$ . E.g., the relation between dynamic theories with and without friction.
- $\mathcal{U}$  is a limiting case of  $\mathcal{V}$ . E.g., the theory of rigid solid objects corresponds to the theory of elastic solid objects in the limit as the elasticity goes to zero. Classical

mechanics corresponds to relativistic mechanics in the limit as the speed of light goes to infinity.

- $\mathcal{U}$  adds more mathematical precision to  $\mathcal{V}$ . E.g., the relation between a theory in which terrestrial gravitation is taken as constant and one in which it is taken as diminishing with elevation.
- $\mathcal{U}$  is a discretized form of  $\mathcal{V}$  obtained by selecting key states of  $\mathcal{V}$  and treating the transitions between these states as atomic actions or events. E.g., the relation between the representation of the blocks world in STRIPS and its representation in solid object dynamics.
- $\mathcal{U}$  is a smoothed form of  $\mathcal{V}$  in which elements or events that are discrete in  $\mathcal{V}$  are replaced by a continuous function governed by a differential equation. E.g., the relation between atomic and continuous models of matter; the use of continuous models of animal population.
- $\mathcal{U}$  and  $\mathcal{V}$  conceptualize the domain in radically different ways. E.g., the relation between wave and particle theories.

It should not be taken for granted that simplifying the form of the theory will make it easier to solve the problem at hand. For instance, problems in statics, where objects are in a stable position and will stand still, are often easier than the same problem in kinematics, where one has to consider all possible motions of the objects that do not make them overlap. Similarly, a theory with friction is simpler to use than a theory without friction in the common case where the friction serves to hold the objects in a fixed position.

*Ignoring small quantities.* For instance, if a problem involving solid objects takes place over regions at different temperatures, it is often reasonable to ignore thermal expansion and contraction, though occasionally, of course, these are critical. Relativistic corrections are ignored in almost all problems that do not involve speeds comparable to the speed of light.

*Dimension reduction.* Dimensions that are irrelevant or along which there is little change may be projected out of a problem. For instance, a problem that involves little change over time may be treated as an atemporal problem. A problem involving moving objects on a surface may be treated as a two dimensional problem. A problem involving moving objects on a track may be treated as a one-dimensional problem. Alternatively, particular entities in the problem may be treated as possessing fewer dimensions than they actually do. For instance, a ball may be abstracted as a point object; a rod in a linkage may be abstracted as a line segment.

Finally, dimension reduction may be carried out in abstract spaces. Consider, for example, a train of  $n$  gears that do not mesh tightly. The coordinated motion of the gears, where they all rotate in sync, constitutes a path through configuration space. More precisely, there is one path through configuration space corresponding to the case where the gear train is moving in one direction, and the front edges of the teeth of the  $k$ th gear meet the back edge of the teeth of the  $(k + 1)$ st; and there is a slightly



different path through configuration space corresponding to the case where the gear train is moving in the opposite direction, and the back edges of the teeth of the  $k$ th gear meet the front edge of the  $(k + 1)$ st. In between these two paths, there is a narrow  $k$ -dimensional tube of configurations, corresponding to the free play of the gears in the small angle range where their teeth do not meet. For many purposes, the radius of the tube can be ignored, and the system can be analyzed as if the configuration space contained only the central path [48].

An extensive survey by Joskowicz and Sacks [36] of the kinematics of 2500 mechanisms in a standard encyclopedia of mechanisms [1] found that some kind of dimensional reduction is possible for the analysis of most mechanisms; it is only a minority of mechanisms that require a full three-dimensional representation of the parts involved.

*Object coalescence.* A collection of objects whose internal relations are fixed can sometimes be treated as a single object. For instance, a table can be treated as a single solid object, rather than reasoning separately about the top, the legs, and the screws as separate interacting objects. (This abstraction breaks down exactly when the table itself breaks down.) A fabric can be treated as a single object made of cloth rather than as a large number of interacting threads.

*Hierarchical analysis of devices.* A complex mechanism can be analyzed as a hierarchy of components at different scales and levels of abstraction. An archetypal, though of course extremely difficult, example is the analysis of an organism as decomposed into organs, tissues, and cells, and sub-cells. This kind of analysis has been carried out with some success for electronic systems [57], but it is in general difficult, first, because it is hard to find a systematic language to characterize the functionality or behavior of high-level components, and second because in order to achieve efficiency, systems are often designed so that high-level modules share sub-parts. The same problems arise in the hierarchical analysis of plans.

Some types of abstraction are easy to carry out computationally but difficult to characterize logically. One such is the abstraction mentioned in Section 14.2.1 in which a kinematic joint is abstracted as a constraint in configuration space. Computationally, such constraints are easily incorporated into the routines that compute the configuration space; once the configuration space has been computed, all subsequent calculations are done purely on the basis of the configuration space and it no longer matters how the configuration space was computed. From a logical point of view, things are more complicated. Are these constraints reified as entities or stated as axioms? If they are reified, then the theory of kinematics must be rewritten to describe the properties of “constraints” and to state how “constraints” enter into the laws of motion. If they are axioms, then there is no longer a clean separation between the problem-specific description of the physical system and the problem-independent physical theory; rather, part of the description of the physical system consists of physical laws (the constraints) that are generated outside the theory itself. Moreover, there will have to be meta-logical rules stating what constraint axioms are reasonable; i.e., can actually be implemented in physical systems. Put it another way: The abstraction of the joints as constraints is simple only under a particular computational

approach: The configuration space is computed from the system description and all further calculations are done from the configuration space, without referring back to the original geometry. But a logical representation does not mandate any particular computational technique, and specifically it cannot prohibit a reasoner from combining results derived from the configuration space with the original system description. This combination will be difficult if there are aspects of the configuration space that are not derived from the original system description.

## 14.4 Historical and Bibliographical

The history of research in AI physical reasoning is punctuated by three major landmarks: in decreasing order of impact, these are

- The publication of the three major qualitative reasoning programs—de Kleer and Brown’s ENVISION [17], Forbus’ QP [23] and Kuipers’ QSIM [39]—in a special issue of *Artificial Intelligence* in 1984. (This was republished as a book a year later [4].) These were in many respects outgrowths of de Kleer’s NEWTON program [16], and Forbus’ FROB [22] which carried out qualitative reasoning for a roller coaster on a track and for balls bouncing among fixed obstacles, respectively. (In both of these programs the moving objects were modeled as point objects.) NEWTON was the first substantial study of commonsense physical reasoning in the AI literature.
- The publication of Pat Hayes’ “Naive Physics Manifesto” [32] and “Ontology for Liquids” [33] in 1978–1979. (The latter circulated for years as a photocopied working paper, until finally being published in 1985.)
- The application of configuration space techniques to problems in solid object kinematics by Faltings [20] and Joskowicz [34] independently in 1987.

Most of the work in physical reasoning relates fairly directly to one of these three.

The very large body of research associated with the qualitative reasoning programs ENVISION, QP, and QSIM is surveyed in Chapter 9, and it would be redundant to repeat that here.

### 14.4.1 Logic-based Representations

In his 1978 paper, “The Naive Physics Manifesto”, Pat Hayes [32] argues the following points. First, an effective strategy in automating commonsense reasoning is to study the logical structure of reasoning in various domains prior to, and largely independently of, considering issues of implementation or application. Second, physical reasoning will be a fruitful domain for this kind of research. Third, commonsense knowledge of physics divides naturally into “clusters” of concepts and axioms, and an effective research strategy will be to axiomatize the clusters separately and then combine the axiomatizations. Fourth, the concept of a “history”, a region in space–time, will be a powerful tool in axiomatizing physical knowledge. Hayes then initiated his research program with “Ontology for Liquids” [33], described above in Section 14.2.3.

“The Naive Physics Manifesto” has inspired and encouraged two separate parts of the KR research community in two different ways. One group of researchers

has embraced the endorsement of research into representations at the logical level, though without being particularly interested in physical reasoning. Another group of researchers has embraced the interest in physical reasoning, but with no enthusiasm about logic. Only a rather small body of work actually attempts to continue Hayes' programme of logical analysis of physical reasoning.

Schmolze [53] presents an axiomatization for a domain that includes actions, events, processes, liquids, solid containers, and faucets. A liquid is modeled as a collection of "granules".

Sandewall [52] developed a logical description of a microworld of points objects moving along surfaces. The chief focus of this work was integrating nonmonotonic logic with a continuous model of time.

Three parallel papers by Lifschitz, Morgenstern, and Shanahan [41, 45, 54] axiomatize various aspects of the process of cracking an egg into a bowl.

Bennett et al. [3] present an axiomatization of solid object kinematics built up from geometrical primitives.

Davis has developed a number of first-order axiomatizations for physical domains, and shown how they can be applied to commonsense inference. An axiomatization of a small part of solid object dynamics, sufficient to support the inference that a marble dropped in a funnel will fall out the bottom is given in [9]. The most significant technical innovation here is the concept of a "pseudo-object", a geometric entity that "moves around" with a rigid object, such as the hole of a doughnut or the center of mass of an object. Chapter 7 of [10] gives preliminary axiomatizations for a number of physical domains, including liquids. An axiomatization of qualitative process theory is given in [11]. The main issue here is to formulate the closed world assumptions correctly.

An axiomatization of a kinematic model of one solid object cutting another is given in [12]. Two theories are presented. The "object" theory views the process of a blade cutting a target object as involving a continuous change in the shape of the target until it splits, when it becomes two objects. The "chunk" theory views the same process in terms of the chunks of solid material contained in the target. (Every separate region defines a separate chunk.) A chunk persists until it is penetrated by the blade, at which point it ceases to exist.

Davis' "Naive Physics Perplex" [15] reconsiders the methodology promoted in Hayes' "Naive Physics Manifesto", and advocates a methodology based around microworlds rather than clusters.

#### 14.4.2 Solid Objects: Kinematics

The idea of configuration space was first developed in robotics to characterize the motions of a robot [42]. Faltings [20] analyzes in detail the kinematics of two-dimensional mechanisms composed of parts each with one degree of freedom, such as mechanical clocks. Juskowicz [34] studies the kinematics of a system that has few degrees of freedom by virtue of the interaction of its components. Forbus et al. [24] carry out a qualitative analysis of a kinematic system, based on the topology of configuration space. Gelsey [28] discusses the construction of kinematic models of varying degree of detail from the geometric specification of a physical system and the use of kinematic models in prediction. Juskowicz and Addanki [35] propose methods for designing the

shape of a kinematic system given a specification of desired properties of the configuration space. Joskowicz and Sacks [36] survey the mechanisms enumerated in a standard encyclopedia of mechanisms and analyzed the complexity of the kinematic analysis required. The robustness of kinematic analysis if it is assumed that shape descriptions are only accurate to within a specified tolerance is discussed in [37] and [13].

### 14.4.3 Solid Object Dynamics

Simulators for the behavior of solid objects using a full dynamic theory have been developed in the contexts of computer-aided engineering [58] and of AI [29]. These carry out an exact simulation of behavior given exact geometric specifications of the objects involved. Sacks and Joskowicz [51] present an algorithm that efficiently carries out dynamic simulation for two-dimensional assemblies using configuration spaces to expedite the problem of collision detection. WHISPER [27] simulates dynamic behavior of two-dimensional systems of solid objects in an occupancy array representation.

The CLOCK program of Forbus, Nielsen, and Faltings [24] extends the qualitative kinematic analysis of [48] with a qualitative representation of forces and motions, thus producing a system for qualitative dynamic prediction. The system takes as input a scanned photograph of a mechanical system such as a mechanical clock with gears, computes the exact configuration space, simplifies and abstracts the configuration space to a qualitative representation, and uses the qualitative configuration space to construct qualitative predictions of behavior. The work of Stahovich, Davis, and Shrobe [55] is similar in spirit to [24]; it is more restricted in scope but more elegant and systematic. This program does qualitative simulation for planar systems of objects, each of which moves with one degree of freedom under the quasi-static assumption that the inertia of objects is negligible as compared to the driving forces and frictional (dissipative) forces, and that collisions are inelastic. The input to the program is a representation of the “qc-space”, which gives, for each pair of interacting objects, a qualitative description of the configuration space of the feasible (non-overlapping) positions and the contact positions of the two objects. (The paper vaguely states that the qc-space can be computed from an informal sketch of the mechanism, but it is not at all clear how this is to be done.) The possible qualitative behaviors of the mechanism is then predicted in terms of trajectories through qc-space, using rules for balancing forces.

### 14.4.4 Abstraction and Multiple Models

The use of multiple models for physical reasoning is proposed in [2]. General studies of the use of abstraction in physical reasoning include [19, 46, 47, 59, 60, 8]. Studies of abstraction in solid object kinematics include [48, 14].

### 14.4.5 Other

Collins and Forbus [7] describe a program that reasons about liquids qualitatively as collections of small particles. The particles are large enough that they can be characterized by thermodynamic properties such as temperature, but small enough that they remain undivided. Gardin and Meltzer [30] simulate rigid objects, flexible objects,

liquids, and strings in terms of interacting molecules. DeCuyper et al. [18] propose a hybrid architecture for reasoning about liquids, combining a qualitative theory similar to Hayes', a particle-based model similar to Gardin and Meltzer's, and a model based on tracking the motion of liquid through a fixed fine-grained decomposition of space.

Rajagopalan [49] uses a qualitative representation of shape and motion to predict magnetic flux and induced current.

Specialized expert systems for specific reasoning in the physical sciences date back to DENDRAL [6], which inferred molecular structure from mass spectroscopy data. But these are highly specific to a narrow domain and task, and hardly connected to more general physical reasoning, either in the knowledge or in the methods of inference used.

An ambitious long-term project, called Project Halo, is underway to encode scientific knowledge in a knowledge base, the Digital Aristotle [25, 26]. The first stage of this project encoded the knowledge in about a chapter's worth of an introductory college chemistry textbook [5]. The project was attempted by three competing knowledge-engineering teams and achieved a fair degree of success; the three systems achieved about the mean human score on questions in the area from the high school AP chemistry test. The subject matter in this first stage—balancing chemical equations and computing acidity of solutions—was chosen specifically to avoid the issues of spatial reasoning and of commonsense reasoning [25].

Great emphasis was placed in Project Halo on carrying out systematic evaluation. The measure is the success rate on answering questions from the relevant section of the advanced placement high school chemistry exam, both in finding the correct answer, and in explaining the answer. The three competing KR teams were presented with a training set of problems, and then their systems were tested on a separate, previously unseen, test set drawn from the same corpus. The grading of the answers was done by an independent set of domain experts. The translation of the English language AP questions into the input formalism was done by the system designers, but overseen by the administrators of Halo.

However, there has been very little analysis or description published of the actual knowledge or representation used. The knowledge bases are available on the Web; see [25]. The current author's examination of the knowledge base created by the Ontoprise group suggests that the representation was very highly geared toward the particular class of problems involved, and avoids even fundamental issues in the area if they do not appear in AP exam questions, as one would expect of a project done under extreme time pressure aiming toward a specified measure of success. For example, the representation does not seem to have any conception of *time*; its representation of an equation like  $2\text{H}_2 + \text{O}_2 \rightarrow 2\text{H}_2\text{O}$  does not allow the inference that first the hydrogen and oxygen is present but not the water, and later the water is present but not the hydrogen and oxygen. Apparently this aspect of chemical equations is taken for granted by the designers of AP tests, and not tested.

#### 14.4.6 Books

There are three major books in the area. *Qualitative Reasoning about Physical Systems* (D. Bobrow, ed., 1985) [4] is a reprint of the 1984 special issue of *Artificial Intelligence*; it includes the original papers on ENVISION, QP, and QSIM. *Readings in*

*Qualitative Reasoning about Physical Systems* (D. Weld and J. de Kleer, eds., 1989) [61] contains essentially all of the important papers in the area published before 1989; it is still the best source for the field. *Qualitative Reasoning: Modeling and Simulation with Incomplete Knowledge* (Kuipers, 1994) [40] presents the QSIM theory and its extensions.

## Bibliography

- [1] I. Artobolevsky. *Mechanisms in Modern Engineering Design*. MIR Publishers, Moscow, 1979.
- [2] S. Addanki, R. Cremonini, and J.S. Penberthy. Reasoning about assumptions in graphs of models. In *Proc. IJCAI-89*, pages 1432–1438.
- [3] B. Bennett, A.G. Cohn, P. Torrini, and S.M. Hazarika. Describing rigid body motions in a qualitative theory of spatial regions. In *AAAI-00*, pages 503–509.
- [4] D. Bobrow, editor. *Qualitative Reasoning about Physical Systems*. MIT Press, Cambridge, MA, 1985.
- [5] T.L. Brown, H.E. LeMay, and B. Bursten. *Chemistry: The Central Science*. Prentice-Hall, 2003.
- [6] B.G. Buchanan, G.L. Sutherland, and E.A. Feigenbaum. Heuristic DENDRAL: A program for generating explanatory hypotheses in organic chemistry. In B. Meltzer, D. Michie, and M. Swann, editors. *Machine Intelligence*, vol. 4, pages 209–254. Edinburgh University Press, 1969.
- [7] J. Collins and K. Forbus. Reasoning about fluids via molecular collections. In *AAAI-87*, pages 590–595.
- [8] B. Choueiry, Y. Iwasaki, and S. McIlraith. Towards a practical theory of reformulation for reasoning about physical systems. *Artificial Intelligence*, 162:145–204, 2005.
- [9] E. Davis. A logical framework for commonsense predictions of solid object behavior. *Int. Journal of AI in Engineering*, 3(3):125–140, 1988.
- [10] E. Davis. *Representations of Commonsense Reasoning*. Morgan Kaufmann, 1990.
- [11] E. Davis. Axiomatizing qualitative process theory. In *KR-92*, pages 177–188.
- [12] E. Davis. The kinematics of cutting solid objects. *Annals of Mathematics and Artificial Intelligence*, 9:253–305, 1993.
- [13] E. Davis. Approximations of shape and configuration space. NYU Computer Science Tech. Report #703, 1995.
- [14] E. Davis. Approximation and abstraction in solid object kinematics. NYU Computer Science Tech. Report #706, 1995.
- [15] E. Davis. The naive physics perplex. *AI Magazine*, 19(4):51–79, Winter 1998.
- [16] J. de Kleer. Multiple representations of knowledge in a mechanics problem solver. In *IJCAI-77*, pages 299–304.
- [17] J. de Kleer and J.S. Brown. A Qualitative physics based on confluences. In D. Bobrow, editor. *Qualitative Reasoning about Physical Systems*, pages 7–83. MIT Press, Cambridge, MA, 1985.
- [18] J. DeCuyper, D. Keymeulen, and L. Steels. A hybrid architecture for modelling liquid behavior. In J. Glasgow, B. Chandrasekaran, and N.H. Narayanan, editors. *Diagrammatic Reasoning*, pages 731–751. MIT Press, 1995.

- [19] B. Falkenhainer. Ideal physical systems. In *Proc. AAAI-93*, pages 600–605.
- [20] B. Faltings. Qualitative kinematics in mechanisms. In *Proc. IJCAI-87*, pages 436–443.
- [21] S. Fahlman. A planning system for robot construction tasks. *Artificial Intelligence*, 4:1–49, 1974.
- [22] K. Forbus. Spatial and qualitative aspects of reasoning about motion. In *Proc. AAAI-80*, pages 170–173.
- [23] K. Forbus. Qualitative process theory. In D. Bobrow, editor. *Qualitative Reasoning about Physical Systems*, pages 85–186. MIT Press, Cambridge, MA, 1985.
- [24] K. Forbus, P. Nielsen, and B. Faltings. Qualitative spatial reasoning: The CLOCK project. *Artificial Intelligence*, 51:417–471, 1991.
- [25] N. Friedland, et al. Toward a quantitative platform-independent quantitative analysis of knowledge systems. In *KR-2004*, pages 507–514.
- [26] N. Friedland. Project halo: toward a digital Aristotle. *AI Magazine*, 25(4):29–47, Winter 2004.
- [27] B. Funt. Problem-solving with diagrammatic representations. *Artificial Intelligence*, 13(3):201–230, 1980.
- [28] A. Gelsey. Automated reasoning about machine geometry and kinematics. In D. Weld and J. de Kleer, editors. *Readings in Qualitative Reasoning about Physical Systems*, pages 580–591. Morgan Kaufmann, 1989.
- [29] A. Gelsey. Automated reasoning about machines. *Artificial Intelligence*, 74:1–53, 1995.
- [30] F. Gardin and B. Meltzer. Analogical representations of naive physics. In J. Glasgow, B. Chandrasekaran, and N.H. Narayanan, editors. *Diagrammatic Reasoning*, pages 669–689. MIT Press, 1995.
- [31] F. Giunchiglia and T. Walsh. A theory of abstraction. *Artificial Intelligence*, 57:323–389, 1992.
- [32] P. Hayes. The naive physics manifesto. In D. Michie, editor. *Expert Systems in the Microelectronic Age*. Edinburgh University Press, Edinburgh, 1979.
- [33] P. Hayes. Naive physics 1: Ontology for liquids. In J. Hobbs and R. Moore, editors. *Formal Theories of the Commonsense World*. Ablex Pubs, Norwood, NJ, 1985.
- [34] L. Joskowicz. Shape and function in mechanical devices. In *IJCAI-87*, pages 611–615.
- [35] L. Joskowicz and S. Addanki. From kinematics to shape: an approach to innovative design. In *Proc. AAAI-88*, pages 347–352.
- [36] L. Joskowicz and E. Sacks. Computational kinematics. *Artificial Intelligence*, 51:381–416, 1991.
- [37] L. Joskowicz, E. Sacks, and V. Srinivasan. Kinematic tolerance analysis. *Computer-Aided Design*, 29(2):147–157, 1997.
- [38] H. Kim. Qualitative reasoning about fluids and mechanics. PhD thesis, Institute of Learning Sciences, Northwestern University, 1993.
- [39] B. Kuipers. Qualitative simulation. In D. Bobrow, editor. *Qualitative Reasoning about Physical Systems*, pages 169–203. MIT Press, Cambridge, MA, 1985.
- [40] B. Kuipers. *Qualitative Reasoning: Modeling and Simulation with Incomplete Knowledge*. MIT Press, 1994.

- [41] V. Lifschitz. Cracking an egg: An exercise in formalizing commonsense reasoning. In *Fourth Symposium on Logical Formalizations of Commonsense Reasoning*.
- [42] T. Lozano-Perez. Spatial planning: A configuration-space approach. *IEEE Transactions on Computers*, C-32:108–120, 1983.
- [43] M. Mason. On the scope of quasi-static pushing. In *Proc. 1985 Third Int. Symp. on Robotics Research*. MIT Press, 1985.
- [44] M. Minsky. A framework for representing knowledge. In J. Haugeland, editor, *Mind Design: Philosophy, Psychology, Artificial Intelligence*, 1981, pages 95–128.
- [45] L. Morgenstern. Mid-sized axiomatizations of commonsense problems: a case study in egg cracking. *Studia Logica*, 67:333–384, 2001.
- [46] P. Nayak. Causal approximations. *Artificial Intelligence*, 70:277–334, 1994.
- [47] P. Nayak and A. Levy. A semantic theory of abstractions. In *IJCAI-95*, pages 196–203, 1994.
- [48] P. Nielsen. A qualitative approach to mechanical constraint. In *Proc. AAAI-88*, pages 270–274.
- [49] R. Rajagopalan. A model for integrated qualitative spatial and dynamic reasoning about physical systems. In *Proc. AAAI-94*, pages 1411–1417.
- [50] E. Sacks and L. Joskowicz. Automated modelling and kinematic simulation of mechanisms. *Computer-Aided Design*, 25(2):106–118, 1993.
- [51] E. Sacks and L. Joskowicz. Dynamical simulation of planar assemblies with changing contacts using configuration spaces. *J. Mechanical Design*, 120, 1998.
- [52] E. Sandewall. Combining logic and differential equations for describing real-world systems. In *KR-89*, pages 412–420.
- [53] J. Schmolze. Physics for robots. In *AAAI-86*, pages 44–50.
- [54] M. Shanahan. An attempt to formalise a non-trivial problem in common sense reasoning. *Artificial Intelligence*, 153:141–165, 2004.
- [55] T. Stahovich, R. Davis, and H. Shrobe. Qualitative rigid-body mechanics. *Artificial Intelligence*, 119:19–60, 2000.
- [56] D. Stewart. Rigid-body dynamics with friction and impact. *SIAM Review*, 41(1):3–39, 2000.
- [57] G. Sussman and G.L. Steele. CONSTRAINTS—A language for expressing almost hierarchical descriptions. *Artificial Intelligence*, 14:1–40, 1980.
- [58] R.A. Wehage and E.J. Haug. Dynamic analysis of mechanical systems with intermittent motion. *J. Mechanical Design*, 104:784–788, 1982.
- [59] D. Weld. Approximation reformulations. In *Proc. AAAI-90*, pages 407–412, 1990.
- [60] D. Weld. Reasoning about model accuracy. *Artificial Intelligence*, 56:255–300, 1992.
- [61] D. Weld and J. de Kleer. *Readings in Qualitative Reasoning about Physical Systems*. Morgan Kaufmann, San Mateo, CA, 1989.