Approximation and Abstraction in Solid Object Kinematics

Ernest Davis*
New York University
251 Mercer St.
New York, New York
DAVISE@CS.NYU.EDU

September 26, 1995

Abstract

Physical reasoning often involves approximating or abstracting the situation or the theory at hand. This paper studies the nature of approximation and abstraction as applied to the kinematic theory of rigid solid objects. Five categories of approximation are considered:

1. Geometric approximation.
2. Abstraction of a complex kinematic structure by a simpler kinematic structure. For example, the abstraction of a collection of tightly-linked objects as a single object.
3. Abstraction of a kinematic structure by a simpler theory. For example, the abstraction by a connectivity graph in configuration space.
4. Approximation of a complex kinematic structure by more complex theory. For example, the approximation of a chain by a string.
5. Approximation of a more complex theory by a kinematic theory. For example, the approximation of solid object dynamics by kinematics.

We discuss how some of these types of approximation can be implemented and combined. We conclude that abstraction and approximation are open-ended styles of reasoning, rather than neatly categorizable meta-relationships.

1 Approximate reasoning

A reasoner encounters a real, physical situation and wishes to answer some question about it. She proceeds by invoking a detailed theory of the physical phenomena at hand and carefully measuring the relevant physical parameters in her situation. She now has a well-defined formal problem $\mathcal{P}$ which is susceptible to calculation. Unfortunately, she finds that it is computationally infeasible to compute the information she wants directly from $\mathcal{P}$. In this position, one option that the reasoner may explore is to consider instead a different, problem $\mathcal{A}$ which may be a less accurate account of the real situation $\mathcal{P}$, but which, she believes, is similar to $\mathcal{P}$ in key respects and easier to compute with. In many situations, the simpler theory $\mathcal{A}$ may also have the advantage of being more robust

*This research was supported by NSF grant #IRI-9300446. Thanks to Leo Joskowicz for many helpful comments and suggestions; also to Dana Angluin, Mladen Covic, Drew McDermott, and Ken Yip.

than $\mathcal{P}$, less dependent on doubtful assumptions or more stable numerically [Gelsey, 95b]. Reasoning and calculation is then performed on $\mathcal{A}$ and the reasoner hopes that the results thus obtained are correct or nearly so.

The relation between $\mathcal{A}$ and $\mathcal{P}$ is is called *approximation* when the dominating differences between $\mathcal{A}$ and $\mathcal{P}$ are in numerical values or other mathematical quantities; *abstraction* when $\mathcal{A}$ is a structural simplification of $\mathcal{P}$; *idealization* when $\mathcal{A}$ is, in some sense, derived by cleaning messy details out of $\mathcal{P}$; or *analogy* if the similarities between $\mathcal{A}$ and $\mathcal{P}$ are weak but suggestive. (Both $\mathcal{A}$ and $\mathcal{P}$ are approximations or idealizations of the actual reality, in some sense; but in what sense is not at all easy to say, and is not discussed here.) In our discussion we will use the term "approximation" to include all these categories. The above use of approximation, where a reasoner starts with a precise theory, and deliberately chooses an approximation to simplify reasoning, will be called "approximation toward simplicity."

A different use of approximation arises in circumstances where the reasoner never has precise or complete information about the physical situation. The sources of the information (perception, experiment, text, inference) may give only partial information; or there may be no very precise theory known for the physical phenomena involved; or the reasoner may be considering a hypothetical situation, such as a system that is under design and only partially specified; or the reasoner may wish to consider a class of situations rather than a single situation. One type of representation for this kind of partial knowledge uses a specific problem $\mathcal{A}$, which is explicitly stated to be an approximation or abstraction of the true problem in specified dimensions to a specified accuracy. For example, one might approximate the shape of an object as a unit cube, and state that this approximation is accurate to within 0.1. This use of approximation, as an expression of partial knowledge, will be called "approximation from ignorance". It is often useful to posit that here, as in approximation towards simplicity, there is a precise formal problem $\mathcal{P}$ that is approximated by $\mathcal{A}$; only here $\mathcal{P}$ is not known.

It is useful to think of a particular approximation as an instance of an *approximation schema.* An approximation schema has three parts. First, there is the class of phenomena being approximated, called the *targets*; second, there is the class of the models being used as approximations, called *nominals*; and, third, there is a *criterion of approximation,* which defines whether a given nominal is a reasonable approximation of a given target. For example, one might choose the target class to be the class of all bounded subsets of the plane; the nominal class to be the class of all finite sets of points; and say that a finite set S approximates a region R within ϵ if every point of R lies within ϵ of a point in S.

There are two central issues common to the use of both approximation toward simplicity and approximation from ignorance in a given domain:

1. What approximation schema or schemas are considered?
2. Given a query Q, a target T, and an nominal I that approximates T relative to a specified criterion, is the answer to Q on T the same or close to the answer to Q on I?

The two uses of approximation differ in the source of the particular nominal. In approximation towards simplicity, the nominal is chosen; in approximation from ignorance, the nominal is given. Therefore, a central issue in approximation towards simplicity is the algorithmic problem

3. How can a suitable nominal be calculated from a given target?

The corresponding issue in approximation from ignorance is the epistemic / representational problem

4. What kinds of partial knowledge are generated by the knowledge source?

Methodologically, (4) is often problematic. It is difficult to characterize the kinds of partial information provided even by a specific knowledge source, such as a particular vision system. But if you want to develop a general-purpose system for reasoning with partial information in a given domain, then what you really want to do is to characterize the partial information provided by any plausible hypothetical knowledge source, which is purely conjectural.

The question of how an AI physical reasoning system could use appropriate idealizations and approximations has received considerable attention recently [Falkenhainer, 93], [Nayak, 94], [Nayak and Levy, 94], [Weld, 90], [Weld, 92]. These studies have primarily focussed on approximation towards simplicity. In general, they have looked at a small number of approximation schemas, such as discarding small terms in equations, or letting particular physical constants appproach a distinguished limit, that apply in many different physical domains.

This paper takes a different angle; so to speak, we take a horizontal rather than a vertical slice of the issue. Following [Joskowicz, 89] we focus attention on a single central physical domain, the kinematic theory of rigid solid models (KRSO), and consider many different approximation schemas that relate to that domain. We are, so to speak, exploring one small region of the graph of models [Addanki et al. 89] in detail. Our primary purpose is to gain a broad view of the nature of abstraction through the study of many example; the paper presents a discursive survey rather than a well-defined technique or tight argument. At the same time, we focus on a single domain to gain some degree of structure and boundedness to our study.

KRSO is an attractive domain for this study in a number of respects. First, the physical aspects of the domain are simple. There is only one type of physical entity, the rigid objects; there are no physical constants or formulas. All the complexity comes from the richness of geometry. Second, the domain is universally familiar. Everyone knows how to push one rigid object with another. Third, KRSO has been studied extensively, both in the AI literature (e.g. [Faltings, 87], [Joskowicz, 87]) and in the computational geometry literature, since it forms the foundations of the "piano movers" problem. (e.g. [Schwartz, Sharir, and Hopcroft, 87]). Abstractions of KRSO have been considered explicitly in such studies as [Joskowicz, 89] and [Nielsen, 88], though this work has been rarely cited in the more recent general work on abstracting physical theories. Finally, applications of KRSO are ubiquitous. Most physical devices contain some components that are rigid solid objects; there are many devices that contain only such components and that work on purely kinematic principles [Joskowicz and Sacks, 91].

We will consider five broad categories of approximation.

I. Approximations of geometric detail. A scenario is approximated by another scenario of the same structure, but with slightly different geometric properties.

II. Abstraction of structure. An arrangement of rigid solid objects is approximated by a structurally different arrangement.

III. Abstraction of configuration space.

IV. Abstraction of KRSO by a different physical theory.

V. Abstraction of a different theory by KRSO.

In (I) and (II) both $\mathcal{P}$ and $\mathcal{A}$ are problems in KRSO. In (III) and (IV), $\mathcal{P}$ is a problem in KRSO and $\mathcal{A}$ is a problem in a different theory. In (V) $\mathcal{A}$ is a problem in KRSO, and $\mathcal{P}$ is a problem in a different theory.

Of the general questions enumerated at the beginning of this section, we are primarily concerned with (1) characterizing types of approximation and abstraction and (2) inferring that the answer to an approximated problem applies to the target problem. Category (2) includes the problem of mechanical tolerance [Joskowicz, Sacks, and Srinivasan, 96], which has the following form: Properties have been established for an ideal system, but mechanical parts can only be manufactured within a certain tolerance of the nominal shape. How large a tolerance can be tolerated to preserve the desired properties?

The above five categories of abstractions are discussed in sections 3 through 7. (Section 3 is a summary of a longer exposition of geometric approximation in KRSO in the companion paper [Davis, 95].) Section 2 gives a brief review of KRSO. Section 8 sketches how an automated reasoner could use these types of approximation. Section 9 compares our results to other work in this area, and presents our conclusions. The casual reader may wish to skip sections 2 and 3, where the material is dense and technical, and instead read sections 4-9, which are comparatively easy going.

2 Kinematics of Rigid Solid Objects

There are three basic physical laws in KRSO:

- The shape of an object does not change.
- Objects move continuously.
- Objects do not overlap.

It is sometimes useful to consider certain objects as *fixed,* and to postulate

- Fixed objects stay in one place.

To facilitate the analysis of the "approximation" and "abstraction" relations within KRSO, and between KRSO and other theories, it will be helpful define precisely the class of shapes that an object can possess, and to define formally various aspects of KRSO.

We begin with some discussion of geometrical terminology and notation. Space is taken to be $\Re^2$ or $\Re^3$, two- or three-dimensional Euclidean space. Points in space will be denoted using bold-face lower-case letters: **p**,**q**, etc. Sets of points, or *regions* are denoted using bold-face upper-case **R**,**O**, etc.

We write "$\mathrm{d}(\mathbf{p},\mathbf{q})$" to denote the Euclidean distance between points **p** and **q**.

Definition 2.1: A *regular region* is a subset of $\Re^n$ that is non-empty, bounded, and equal to the closure of its interior.

Definition 2.2: Two regions *overlap* if their interiors have a non-empty intersection.

Definition 2.3: A *placement* is a pair $\langle \mathbf{R}, M \rangle$ of a regular region **R** and a rigid mapping M. Intuitively, **R** is the region occupied by an object in some standard position, and M is the displacement of the object from that standard. Thus, in placement $\langle \mathbf{R}, M \rangle$ the object occupies position $M(\mathbf{R})$.

We will indicate the ith component of a tuple V as $V[i]$, reserving subscripts to distinguish different tuples.

Definition 2.4: A *display* is a k-tuple of regular regions. Intuitively, these are the shapes of k objects. Hence, we will often refer to the indices $1 \ldots k$ as "objects".

Definition 2.5: A display D' is a *contraction* of display D if, for $i = 1 \ldots k$, $D'[i] \subset D[i]$. D' is an *expansion* of D if D is a contraction of D'.

Definition 2.6: A *configuration* is a k-tuple of rigid mappings C. Intuitively, these are the displacements of each object from its standard position as given in a display. The *configuration space on k objects* is the set of all such k-tuples.

Definition 2.7: A *scenario* is a pair of a display and a scenario, or, equivalently, a tuple of placements. If $\langle D, C\rangle$ is a scenario, then, slightly abusing notation, we will write $CD[i]$ for $C[i](D[i])$, the region occupied by the ith object in the scenario. $CD[i]$ is called the *place* of object i in scenario $\langle D, C\rangle$.

Definition 2.8: A scenario $\langle D, C\rangle$ is *feasible* if for all $I \neq J$, $CD[I]$ does not overlap $CD[J]$; else it is *forbidden*. $\langle D, C\rangle$ is *contact-free* if for all $I \neq J$, $CD[I] \cap CD[J] = \emptyset$. For any display D, the set of configurations C such that $\langle D, C\rangle$ is feasible is denoted "free(D)"; the set of configurations C such that $\langle D, C\rangle$ is forbidden is denoted "forbidden(D)"; and the set of configurations C such that $\langle D, C\rangle$ is contact-free is denoted "cfree(D)".

Definition 2.9: A *path* is a continuous function from the real interval [0,1] to configuration space. A *history* is a pair $\langle D, \phi\rangle$ where D is a display and ϕ is a path. History $\langle D, \phi\rangle$ is feasible/contact-free if scenario $\langle D, \phi(T)\rangle$ is feasible/contact-free for every T. In cases where we are considering a fixed display, we will use the terms "path" and "history" interchangably, using whichever reads more naturally.

Definition 2.10: Let $\langle D, C\rangle$ be a feasible scenario. The path-connected component of free(D) is denoted "accessible(C, D)".

Many different kinds of queries can be posed within this theory. Probably the most extensively studied is the

Motion planning problem: Given a display D and configurations C_1 and C_2 that are feasible over D, find a path, or find an optimal path, from C_1 to C_2 through free(D), or determine that no such path exists.

Other problems include

Configuration space calculation: Given a scenario $\langle D, C\rangle$, compute or characterize accessible(C, D).

Response to input motion: Given a set of object $\mathcal{O}$, a scenario $\langle D, C\rangle$ over $\mathcal{O}$, and a history I of one object $O \in \mathcal{O}$ (the object being "driven"), compute or characterize all feasible histories over all the objects in $\mathcal{O}$ which begin in C and in which O follows I or some initial segment of I. (If I leads to a "jammed" position, then it may not be possible to follow more than an initial segment of I.)

Design from a configuration space: Given a connected set of configurations $\mathcal{C}$, find a display D such that $\mathcal{C}$ is a path-connected component of free(D).

Design from input/output: Given an input path ϕ_1 and an output path ϕ_2 through configuration space, find a display D and an initial configuration C such that moving object O_1 along path ϕ_1 starting in C will result in object O_2 moving along path ϕ_2.

Motion planning with dynamic constraints: Find a path from scenario $S1$ to scenario $S2$ of duration no more than T that obeys specified constraints on velocity, acceleration, angular velocity, and so on.

In this paper we will focus primarily on motion planning. However, many of our results that characterize the relation between approximations of shape and position and approximations of the configuration space are applicable to any problem in KRSO.

3 Geometric Approximations

In general, the geometric aspects of a kinematic system — shapes, placement, paths, configurations — are not given in full exact detail. Each of these takes on a value from a space which is in principle uncountably infinite and thus not representable in a countable language, and in practice very large and not effectively expressible in a reasonable language. For example, suppose that a cavity one millimeter cubed is a perceptible defect in an object one meter cubed, which for many purposes is true. Then there are considerably more than 2^{10^8} different shapes that are perfect meter cubes except for isolated millimeter cavities. An exact representation would thus require more than 10^8 bits for practically every such shape. For effective representation and computation, therefore, shapes and other geometric entities must be partially characterized; and approximation is the most common technique for this partial characterization. That is, we describe any such object as approximated by a meter cube, to within a millimeter. (To be precise, the true shape is half a millimeter from the nominal shape in the Hausdorff distance, defined in definition 3.1 below.)

The simplest type of approximation is in terms of containment: The nominal region $\mathbf{A}$ may be stated to be a subset or a superset of the target region $\mathbf{T}$. The application of this approximation to motion planning is evident: one can find a solution to the path-finding problem by approximating each shape by a superset and solving the problem on the supersets; one can establish that no solution exists by approximating each shape by a subset and showing that no solution exists on the subset. ([Fleischer et al., 92] use this technique to find efficiently a solution to the path-finding problem in cases where the answer is "easily seen"; they approximate an arbitrary convex shape by an inscribed and a circumscribed triangle. See also [Rajagopalan, 94].)

Theorem 3.1: Let D be a display and let D' be an expansion of D. If path ϕ is physically feasible over D' then it is physically feasible over D.

In [Davis, 95] we study a variety of approximation criteria on shapes, on regions in configuration space, and on sets of feasible paths, and the relation between these. It is not in general true that small changes in shape result in small changes in free space or in the class of free paths; figures 3-6 show some important examples where small changes in shape cause radical changes in behavior. However, it is possible to give some conditions that do suffice to guarantee continuity of behavior under small changes in shape. We cite the major results here; for details of definitions and proofs, see [Davis, 95].

First, we define three notions of distance between shapes

Definition 3.1: (Figure 1) The distance from a point $\mathbf{p}$ to a region $\mathbf{O}$, $\mathrm{d}(\mathbf{p},\mathbf{O})$, is, as usual, the minimum over $\mathbf{q} \in \mathbf{O}$ of $\mathrm{d}(\mathbf{p},\mathbf{q})$. The *Hausdorff distance* between regions $\mathbf{R}$ and $\mathbf{S}$ is the maximum of

i. the maximum over $\mathbf{r} \in \mathbf{R}$ of $\mathrm{d}(\mathbf{r},\mathbf{S})$;

ii. the maximum over $\mathbf{s} \in \mathbf{S}$ of $\mathrm{d}(\mathbf{s},\mathbf{R})$;

Definition 3.2: (Figure 2 The *dual-Hausdorff distance* between regions $\mathbf{R}$ and $\mathbf{S}$ is the maximum of

i. the Hausdorff distance between $\mathbf{R}$ and $\mathbf{S}$

ii. the Hausdorff distance between the complement of $\mathbf{R}$ and the complement of $\mathbf{S}$

Definition 3.3: Let $\mathbf{A}$ and $\mathbf{B}$ be regions with piecewise smooth boundaries; let $\epsilon > 0$ be a length and let $\phi > 0$ be an angle. $\mathbf{B}$ *approximates* $\mathbf{A}$ *in tangent*, with parameters ϵ, ϕ, if there exists a homeomorphism Γ from $\mathbf{A}$ to $\mathbf{B}$ such that,

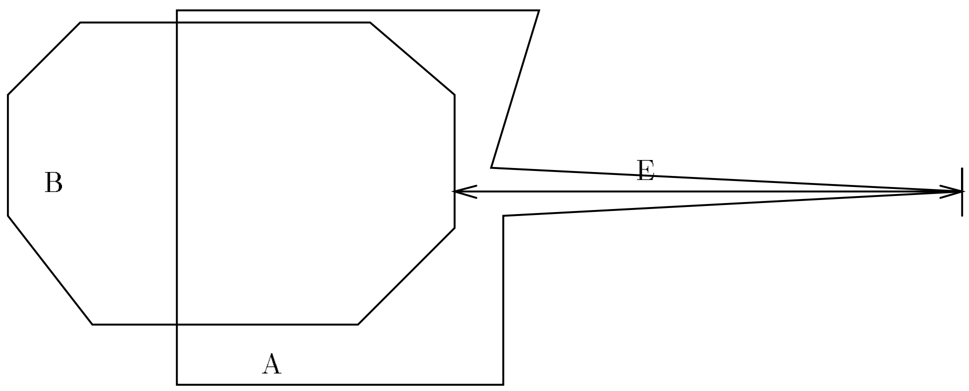

E is the Hausdorff distance between regions A and B.

Figure 1: Hausdorff distance

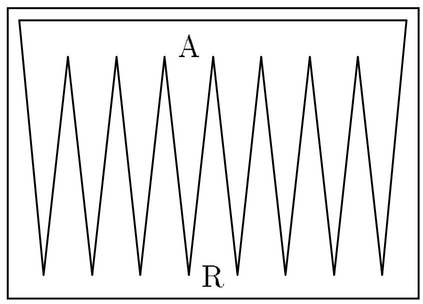

The Hausdorff distance from **R** to **A** is roughly half the width of a tooth.
The dual-Hausdorff distance from **R** to **A** is half the height of **R**.

Figure 2: Dual-Hausdorff distance

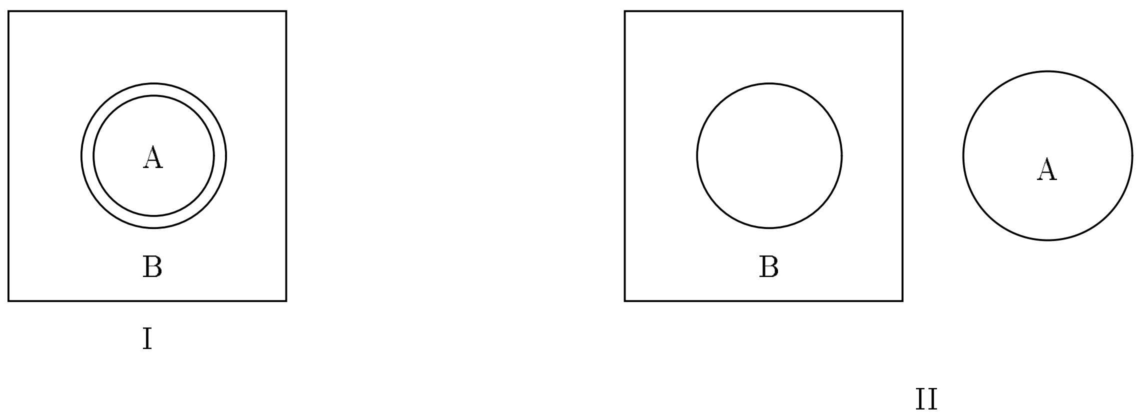

In I, there is a configuration where A is inside B.
This is not close to any feasible configuration in II.

Figure 3: Change in shape vs. change in free space

i. For any $\mathbf{a} \in \mathbf{A}$, $\mathrm{d}(\mathbf{a}, \Gamma(\mathbf{a})) < \epsilon$; and

ii. Let $\mathbf{a} \in \mathrm{Bd}(\mathbf{A})$. Then any tangent to $\mathbf{A}$ at $\mathbf{a}$ lies within ϕ of some tangent to $\mathbf{B}$ at $\Gamma(\mathbf{a})$ and conversely.

We also distinguish three criteria for measuring similarity in configuration space. Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be two sets of configurations (usually the free spaces or contact-free spaces of two displays). Then $\mathcal{C}_1$ can be considered close to $\mathcal{C}_2$ if

I. The dual-Hausdorff distance between $\mathcal{C}_1$ and $\mathcal{C}_2$ is small. (This distance is derived from a natural metric on scenarios discussed in [Davis, 95].)

II. Every path through $\mathcal{C}_1$ is close to a path through $\mathcal{C}_2$ and vice versa.

III. Every path through $\mathcal{C}_1$ can be tracked closely by a path through $\mathcal{C}_2$, and vice-versa.

IV. For each configuration C_1 in $\mathcal{C}_1$ there is a configuration C_2 in $\mathcal{C}_2$ such that every path starting in C_1 can be tracked closely by a path starting in C_2 and vice versa.

Figures 3-6 show examples where change of the object shape is small but (I) is not satisfied; where (I) is satisfied but not (II); where (II) is satisfied but not (III); where (III) is satisfied but not (IV).

We can now establish the following results:

Theorem 3.2: For any display D, if D' is a contraction of D and the dual-Hausdorff distance from D to D' is sufficiently small, then the dual-Hausdorff distance from free(D) to free(D') will be small.

Theorem 3.3: For any display D, if D' is an expansion of D and the dual-Hausdorff distance from D to D' is sufficiently small, then the dual-Hausdorff distance from cfree(D) to cfree(D') will be small.

Theorem 3.4: For any display D, if (i) D' is a contraction of D; (ii) the dual-Hausdorff distance from D to D' is sufficiently small; and (iii) free(D) is "ordinarily connected"; then any feasible path through free(D') can be tracked very closely by a path through free(D) and vice versa.

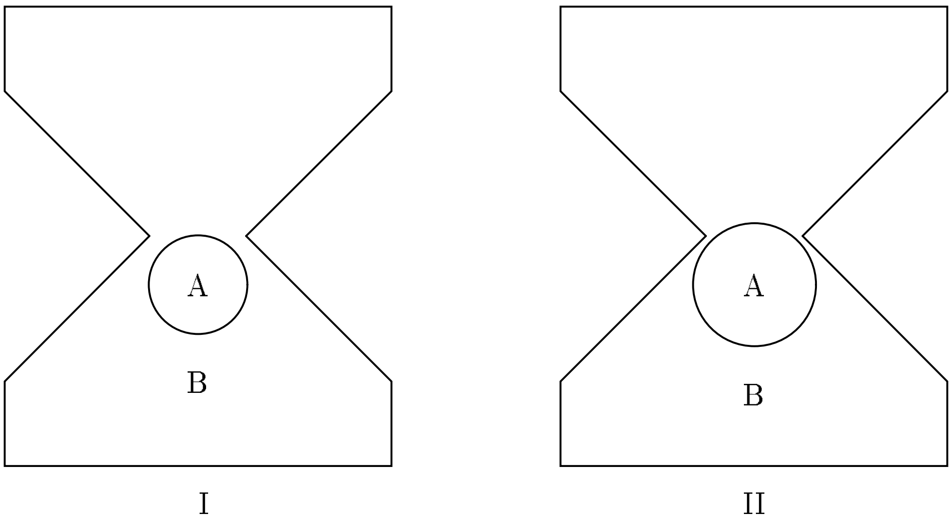

In I, the configurations where the ball is on top and those where it is on the bottom are in a single connected component of free space. In II, they are two separate connected components. Thus, no free path is II is close to the path from bottom to top, which is free in I.

Figure 4: Free space vs. path traces

"Ordinarily connected" means that, in every spherical neighborhood U, every connected component of free(D) $\cap$ U is path-connected. This property is satisfied by any physically reasonable object shapes.

Theorem 3.5: For any display D, if (i) D' is an expansion of D; (ii) the dual-Hausdorff distance from D to D' is sufficiently small; and (iii) free(D) is "ordinarily connected"; then any feasible path through cfree(D) can be tracked very closely by a path through cfree(D') and vice versa.

Theorem 3.6: Let D be a display of piecewise smooth two-dimensional objects. Let D be "always strongly separable"; that is, in any feasible configuration over D, there is a motion that move every two points of two objects in contact "away" from one another. Then if D' approximates D in tangent (ϵ, ψ) for sufficiently small ϵ, ψ, then condition (IV) above is satisfied for the spaces $\mathcal{C}_1$=free(D) and $\mathcal{C}_2$=free(D')

The precise definition of "always strongly separable" in theorem 3.6 is too complex to be given here. Figure 3 shows a case where the display is not always strongly separable, and Theorem 3.6 fails. [Davis, 95] gives formal definitions, proofs, and further related theorems and conjectures.

3.1 Gap closing at the geometric level

One of the most common and important categories of approximation involves the analysis of a system which is slightly loose, and abstracting it as fitting tightly. Joskowicz [1988] calls this "gap-closing". This approximation can serve to reduce dramatically the dimensionality and complexity of the configuration space. (See section 5.4 for an extreme example; also sections 5.2 and 8.) Unfortunately, I have not been able to find any general formal characterizations of this type of approximation. Below is a simple specific example:

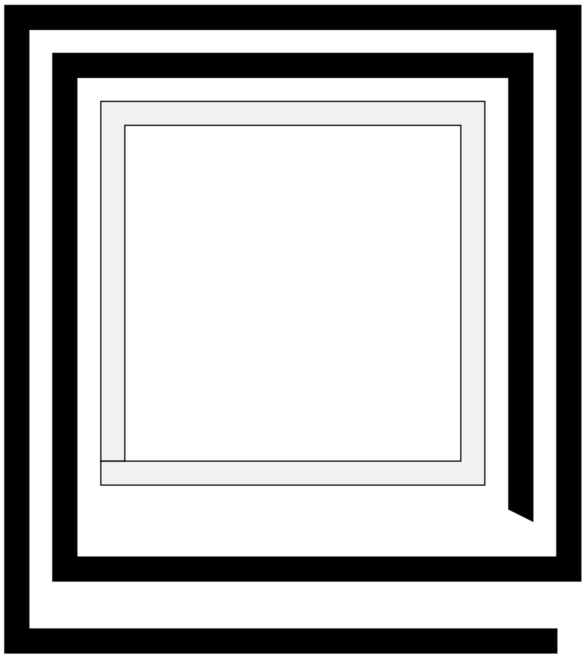

A is the closed square (in grey); B is the spiral (in black)
Any path in A is close to a path in B,
but paths in A that go around more than twice cannot be closely tracked in B.

Figure 5: Path traces vs. tracking paths

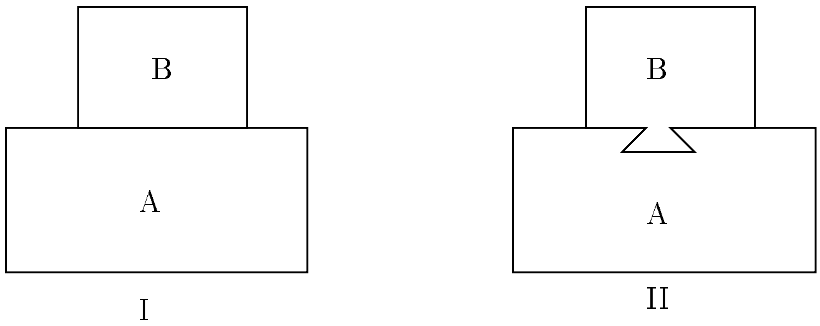

Figure 6: Path tracking vs. close correspondence of points

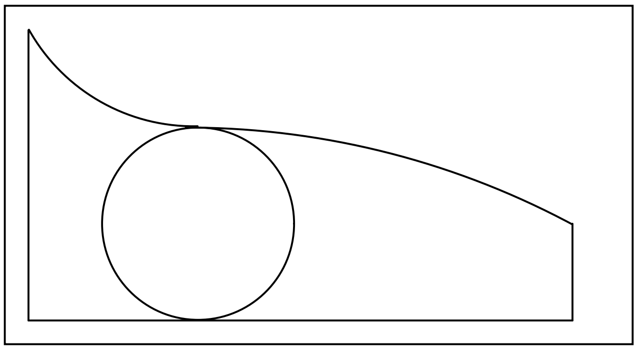

A

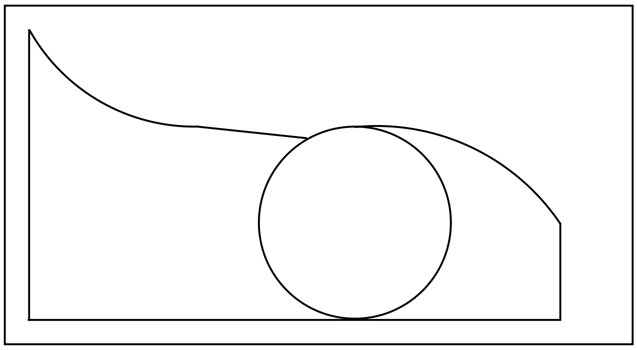

B

In (A) the ball is free to move to the left.
However, a “trapped” position can be created by carving out a “notch” as in (B),
of arbitrarily small depth and change in angle.

Figure 7: Change in connectedness under small change in tangent

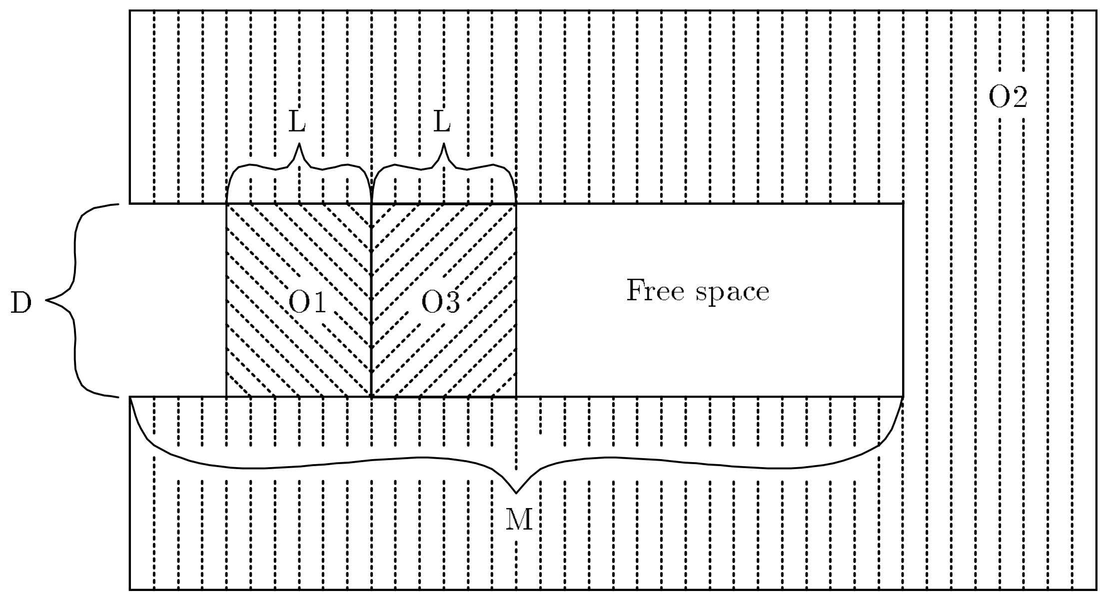

Figure 8: A driven system

Let $\langle D', C\rangle$ be the scenario shown in figure 8 with objects O_1, O_2, O_3. Consider the driving history H starting in C in which O_1 moves horizontally a distance $M - 2L$ and O_2 is fixed. Clearly, in the permitted responses space, O_3 moves horizontally to the right of O_1 and never jams. (Note that, since our theory is purely kinematic, we do not make the "quasi-static" assumption [Mason, 85] that O_3 moves only as far as it is pushed.)

Now, suppose that D' is in fact the nominal approximation of the real display D, which satisfies the following properties:

- The shape of O_1 is the same in D as in D'.
- The true shapes of O_2, O_3 in D are subsets of the nominal shapes in D
- The nominal shapes of O_2, O_3 in D' approximate the true shapes in D in tangent with parameters (ϵ, ϕ).
- The following inequalities hold:

$$\epsilon < D/3 \tag{1}$$

$$\epsilon < L/2 \tag{2}$$

$$\epsilon \leq D + L/2 - \sqrt{D^2 + DL} \tag{3}$$

$$\phi < \pi/2 \tag{4}$$

$$\cos(\phi) \geq \frac{\epsilon}{L - 2\epsilon} \tag{5}$$

$$\tan(\phi) \leq \frac{L - 2\epsilon}{D + 2\epsilon} \frac{(D - 2\epsilon)(D + 2\epsilon) + (L - 2\epsilon)\sqrt{(D - 2\epsilon)^2 + (L - 2\epsilon)^2 - (D + 2\epsilon)^2}}{(D - 2\epsilon)^2 + (L - 2\epsilon)^2} \tag{6}$$

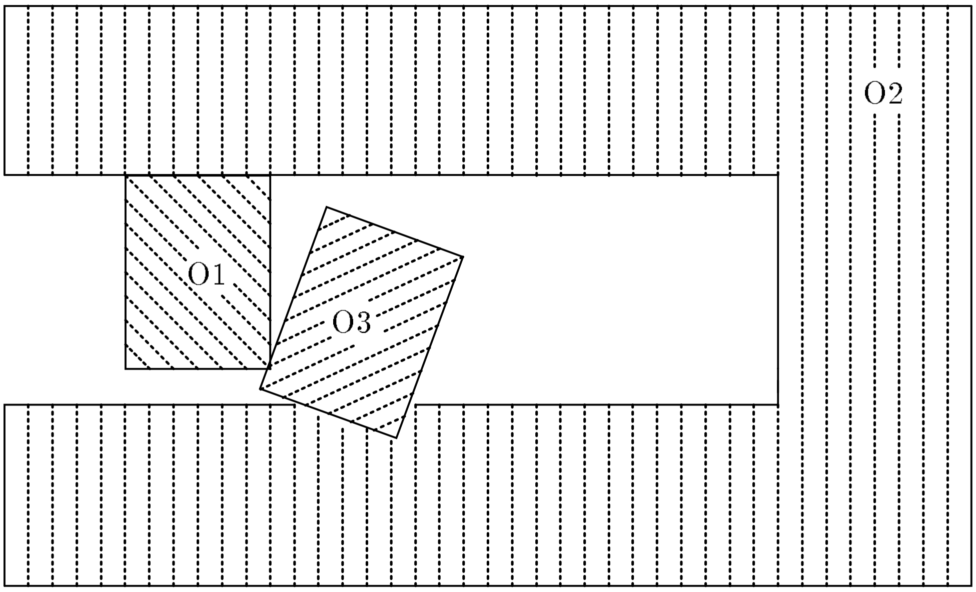

Figure 9: Stuck on one side

Instead of inequality (6) above, the following simpler constraint can be substituted, yielding a sufficient though not necessary set of conditions:

$$\tan(\phi) \leq \frac{L-2\epsilon}{D+2\epsilon}\left(1-\frac{8\epsilon^2}{(D-2\epsilon)^2+(L-2\epsilon)^2}\right) \tag{7}$$

The above inequalities serve the following purposes. (1) ensures that object O_3 does not slip leftward past O_1, passing it above or below. (2) ensures that O_3 contains a rectangle of some width. (3) ensures that O_3 does not have room to rotate freely in the channel. (4) ensures that O_3 cannot jam in a side-to-side contact with the channel. (5) and (6) ensure that the system cannot jam in a position such as that shown in figures 9 and 10 respectively; in any such position, object O_3 will continue to move rightward by having the rightmost edge rotate forward and toward the center.

We then have the following two results:

I. The driven system cannot jam.

II. Any path through accessible(C, D') can be tracked by a path through accessible(C, D) with a separation less than

$$\epsilon(2+4\max(1, \frac{D-2\epsilon}{L-2\epsilon}))$$

and vice versa.

The proof of (I) essentially follows the explanations given above for the inequalities. The proof of (II) proceeds by showing that the path distance is equal to the Hausdorff distance between the free spaces, and then finding the configuration in the target space which is furthest from any configuration in the nominal space. Formalizing these proofs would involve a fairly long and very tedious case analysis.

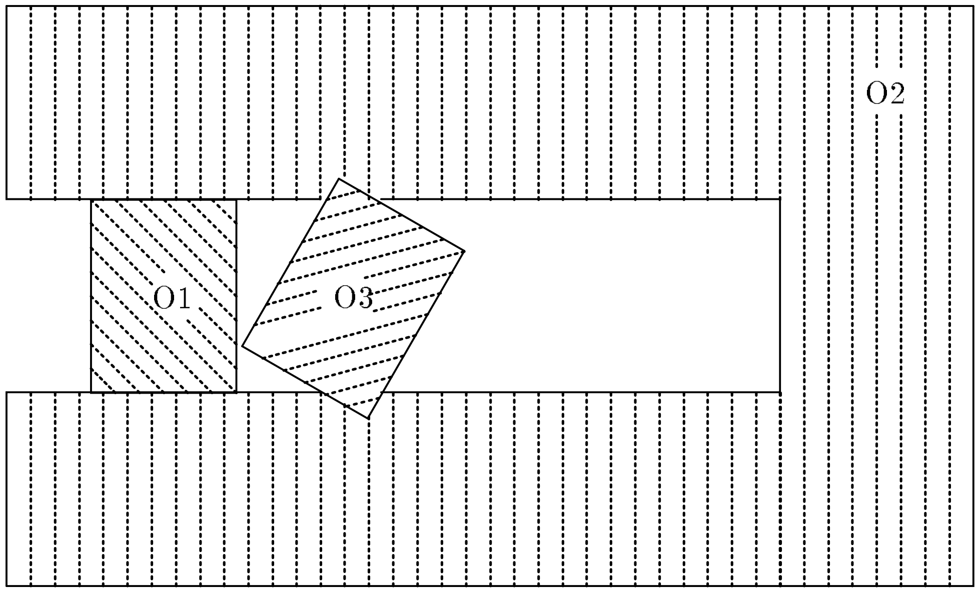

Figure 10: Stuck on two sides

The details of the above inequalities are not particularly important as, in practice, systems are built with tolerances very much less than those allowed by these inequalities. What is important is that some kind of approximation in tangent together with a contraction condition is sufficient to guarantee qualitative similarity in response. I conjecture, but have not proven, that under conditions similar to those of Theorem 3.6, a sufficiently accurate approximation in tangent of the pieces of a mechanism leads to a qualitatively correct prediction of its response to an input motion.

3.2 Irrelevance of the inaccessible

Another technique for approximating shapes uses the fact that parts of an object's boundary that cannot come into contact with any other object are necesarily irrelevant to the kinematic behavior. (Figure 11)

The rule seems at first sight to require circular reasoning; one wishes to approximate shapes in order to simplify the calculation of configuration space, but this rule requires knowing properties of the configuration space in order to justify the shape approximation. This rule finds its use, therefore, in cases where some inaccessibility relations can be calculated quickly, then the shapes are approximated, then the configuration space is calculated in detail.

Theorem 3.7: If no object ever overlaps a region $\mathbf{R}$ attached to object O, then the shape of O can be augmented by $\mathbf{R}$ without changing the accessible configuration space.

Formally: Let $\langle D, C\rangle$ be a scenario and let $\mathcal{C} = \text{accessible}(C, D)$. Let $\mathbf{R}$ be a compact region and O be an object with the following properties:

- $D[O] \cup \mathbf{R}$ is connected.
- For any configuration $C' \in \mathcal{C}$ and $O_1 \neq O$, $C'D[O_1]$ does not overlap $C'[O](\mathbf{R})$.

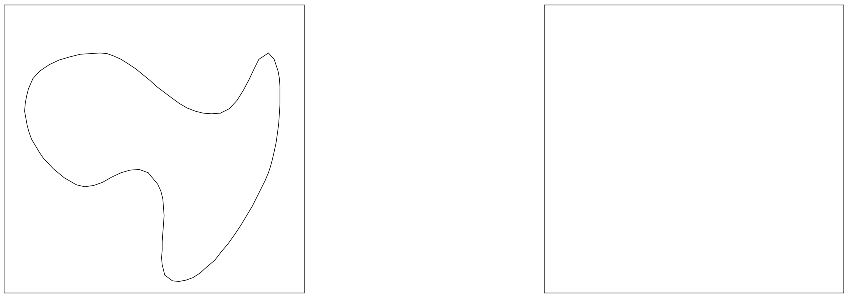

The internal cavity can be ignored.

Figure 11: Abstracting away inaccessible boundary parts

Let D' be the display such that $D'[O] = D[O] \cup \mathbf{R}$; $D'[O_1] = D[O_1]$ for $O_1 \neq O$. Then accessible(C, D') = accessible(C, D).

Proof: Immediate. ▮

Theorem 3.8: If O is an object and $\mathbf{R}$ is a chunk of O such that no other object ever abuts the boundary of $\mathbf{R}$, then $\mathbf{R}$ can be carved out of O without affecting the configuration space.

Formally: Let $\langle D, C \rangle$ be a scenario and let $\mathcal{C}$ = accessible(C, D). Let $\mathbf{R}$ be a compact region and O be an object with the following properties:

- $D[O] - \mathbf{R}$ is non-empty and connected.
- For any configuration $C' \in \mathcal{C}$ and $O_1 \neq O$, $C'D[O_1]$ does not intersect $C'[O](\mathbf{R})$.

Let D' be the display such that $D'[O] =$ closure$(D[O] - \mathbf{R})$; $D'[O_1] = D[O_1]$ for $O_1 \neq O$. Then accessible(C, D') = accessible(C, D).

Proof: Immediate. ▮

In particular, internal cavities — that is, closed subsets of a region that do not intersect the boundary — may be carved out or filled in without affecting the kinematic relations of the object with the outside world.

3.3 Boundary-based approximation criteria

Another shape approximation criterion that has been studied specifies that the boundary of the approximation lies close to the boundary of the target [Joskowicz,Sacks and Srinivasan, 96]. This can be taken in a number of senses.

One possibility is to require that the Hausdorff distance between the boundaries be small. This requirement does not suffice to guarantee that the configuration spaces are nearly equal. (Figure 12)

A second possibility, suggested by Requicha [1983] is this: Define the ϵ-shell of a region $\mathbf{R}$, denoted ϵ-shell$(\mathbf{R})$, to be the set of all points within ϵ of the boundary of $\mathbf{R}$. (Requicha's representation

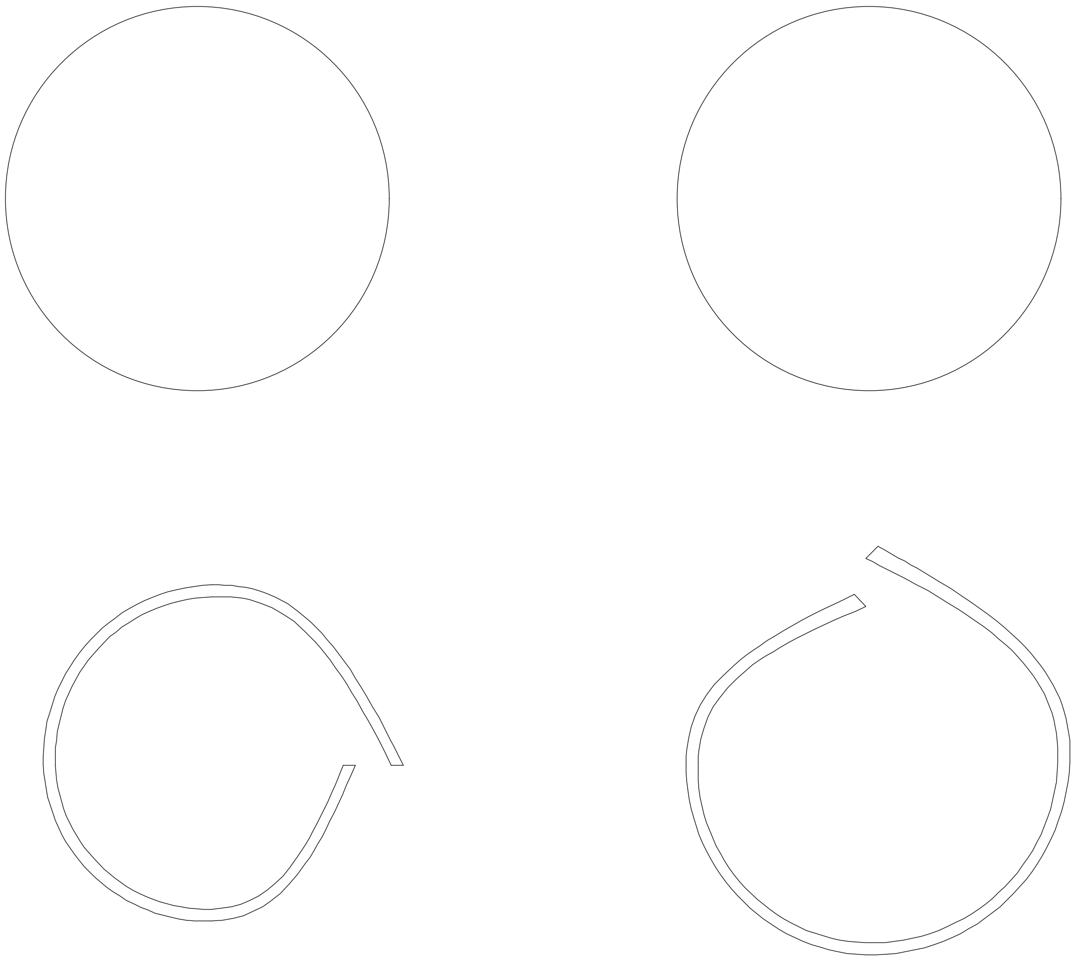

The boundary of a broken bent ring is close to the boundary of a disk, but kinematically they are very different. The smaller ring can be slid inside the larger one, while such a motion of the rings is solidly blocked.

Figure 12: Hausdorff distance between boundaries

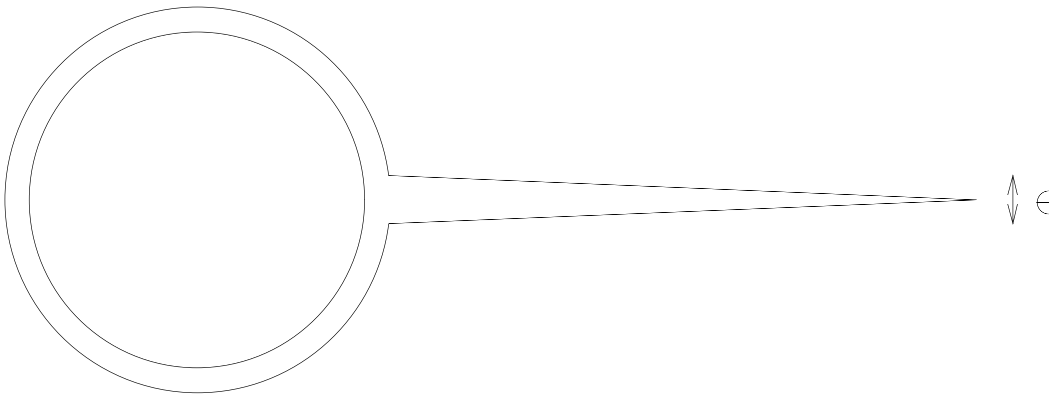

The inner circle is an ϵ-shell approximation of the outer circle with beak.

Figure 13: ϵ-shell approximation

actually allows the possibility of having different tolerances on different parts of the boundary; we ignore this for simplicity.) We can then say that $\mathbf{R}$ is an ϵ-shell approximation to region $\mathbf{S}$ if

$$\mathbf{R} - \epsilon\text{-shell}(\mathbf{R}) \subset \mathbf{S} \subset \mathbf{R} \cap \epsilon\text{-shell}(\mathbf{R})$$

This definition is assymmetric in $\mathbf{R}$ and $\mathbf{S}$. A potential problem with the definition is that features of $\mathbf{R}$ that are thinner than ϵ can vanish altogether (figure 12). It is useful only in the case where the approximation is given as a shape with no thin parts, and specified to be a ϵ-shell approximation of the target shape. (Requicha, being interested in manufacturing tolerances, uses "approximation" in exactly the reverse sense; the manufactured object approximates the nominal shape.)

One way to eliminate this problem is to use the symmetric relation that $\mathbf{R}$ is an ϵ-shell of $\mathbf{S}$ and vice versa. This symmetric relation is equivalent to the condition that both the Hausdorff and the complement-Hausdorff distance from $\mathbf{R}$ to $\mathbf{S}$ are less than ϵ.

4 Approximation of structure within KRSO

In this section, we consider various ways in which one KRSO problem can be abstracted as a different, structurally simpler KRSO problem, either by eliminating objects, coalescing objects, or by reducing the dimensionality of the geometric space.

4.1 Separation of inaccessible objects

If some of the objects in a configuration cannot come into contact with others, then in reasoning about the former one can ignore the latter. In kinematic theories, there is no interaction at a distance. (Figure 14)

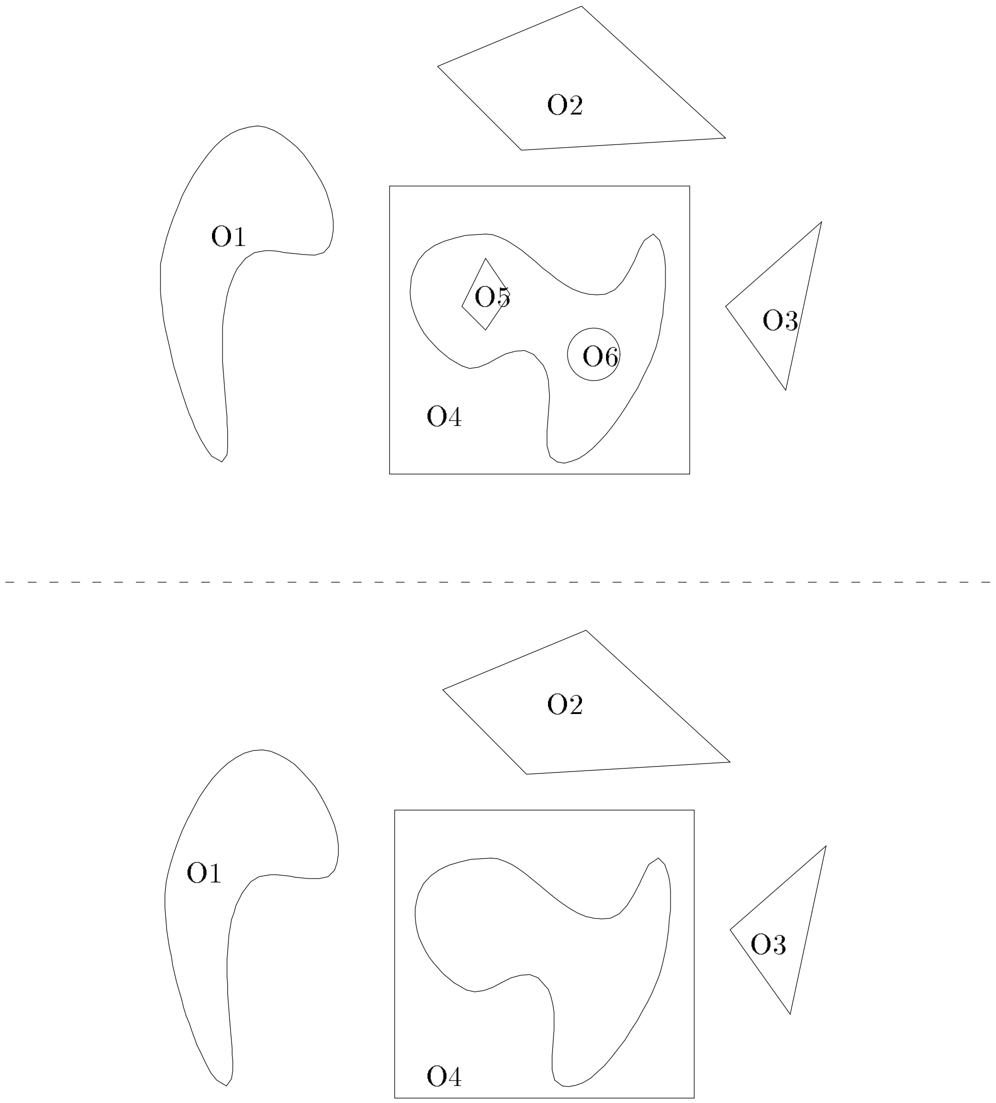

In solving a problem that only refers to objects o1, o2, o3, and o4, it is possible to ignore o5 and o6.

Figure 14: Ignoring inaccessible objects

Theorem 4.1: Let $\mathcal{O}$ be a set of objects, partitioned into three subsets $\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3$. Let $\langle D, C \rangle$ be a scenario over $\mathcal{O}$ and let $\mathcal{C}$=accessible(C, D). Assume that no configuration in $\mathcal{C}$ brings any object in $\mathcal{O}_2$ into contact with an object in $\mathcal{O}_3$. Let $\langle D_{12}, C_{12} \rangle$ be the restriction of $\langle D, C \rangle$ to $\mathcal{O}_1 \cup \mathcal{O}_2$. Then the projection of $\mathcal{C}$ onto $\mathcal{O}_{12}$ is equal to accessible(C_{12}, D_{12}).

Proof: Immediate. ▮

4.2 Elimination of removable objects

If the objects are divided into two parts separated by a plane, then one can solve path-finding problems for the objects on one side independently, ignoring those on the other side. (Figure 15)

Theorem 4.2: Let the set of objects be partioned into two subsets $\mathcal{O}_1$ and $\mathcal{O}_2$. Let $\langle D, C \rangle$ be a feasible scenario and let $\mathbf{R}$ be a plane such that in $\langle D, C \rangle$ the place of every object in $\mathcal{O}_1$ is on one side of $\mathbf{R}$ and the place of every object in $\mathcal{O}_2$ is on the other. Let $\langle D_1, C_1 \rangle$ be the restriction of $\langle D, C \rangle$ to $\mathcal{O}_1$ and let ϕ_1 be a path starting in C_1 through free(D_1). Then there is a ϕ through free(D) such that the restriction of ϕ to $\mathcal{O}_1$ is ϕ_1.

Proof: Let $\hat{z}$ be the vector perpendicular to $\mathbf{R}$, pointing toward the side of $\mathcal{O}_2$. Let $\mathbf{o}$ be a reference point (the origin). Let $z_m(T)$ be the maximum value of $(\mathbf{p} - \mathbf{o}) \cdot \hat{z}$ where $\mathbf{p}$ ranges over points inside the places of the objects in $\mathcal{O}_1$ in situation $\phi_1(T)$. Let $z_M(T)$ be the maximum value of $z_m(t) - z_m(0)$ for $t \in [0, T]$.

Now consider the history $\phi(T)$ where $\phi(T, O) = \phi_1(T, O)$ for $O \in \mathcal{O}_1$ and $\phi(T, O) = CD[O] + z_M(T)\hat{z}$ for $O \in \mathcal{O}_2$. The continuity of ϕ is immediate. The feasibility of ϕ follows from the facts that

- The objects in $\mathcal{O}_1$ undergo the same motions in ϕ as in ϕ_1. Hence, no two objects in $\mathcal{O}_1$ ever overlap.
- The position of the objects in $\mathcal{O}_2$ at time T is a translation by $z_M(T)\hat{z}$ of their position in C. Since they do not overlap in C, they do not overlap in $\phi(T)$.
- In configuration $\phi(T)$, the objects in $\mathcal{O}_2$ are separated from those in $\mathcal{O}_1$ by the plane $\mathbf{R} + z_M(T)\hat{z}$.

▮

4.3 Abstraction of tightly bound objects as a single object

A collection of tightly bound objects can be treated as a single object. (Figure 16)

Definition 4.1: Let $\langle D, C \rangle$ be a scenario. $\langle D, C \rangle$ is *tightly bound* if every configuration in accessible(C, D) differs from C by a single rigid motion applied to all the objects.

Definition 4.2: Let $\langle D, C \rangle$ be a scenario and let $\mathcal{O}$ be a subset of the objects of C. The objects of $\mathcal{O}$ are *tightly bound in* $\langle D, C \rangle$ if $\langle D, C \rangle$ restricted to $\mathcal{O}$ is tightly bound.

Lemma 4.1: If $\mathcal{O}$ is tightly bound in $\langle D, C \rangle$, then the union over $O \in \mathcal{O}$ of $DC[O]$ is a connected region.

Proof: Suppose that the union of the places of the objects is not connected. Then it can be expressed as $\mathbf{R}_1 \cup \mathbf{R}_2$, where $\mathbf{R}_1$ and $\mathbf{R}_2$ are disjoint and closed. Clearly, no object can occupy a region that intersects both $\mathbf{R}_1$ and $\mathbf{R}_2$, so let $\mathcal{O}_1$ be the set of objects whose place is a subset of $\mathbf{R}_1$

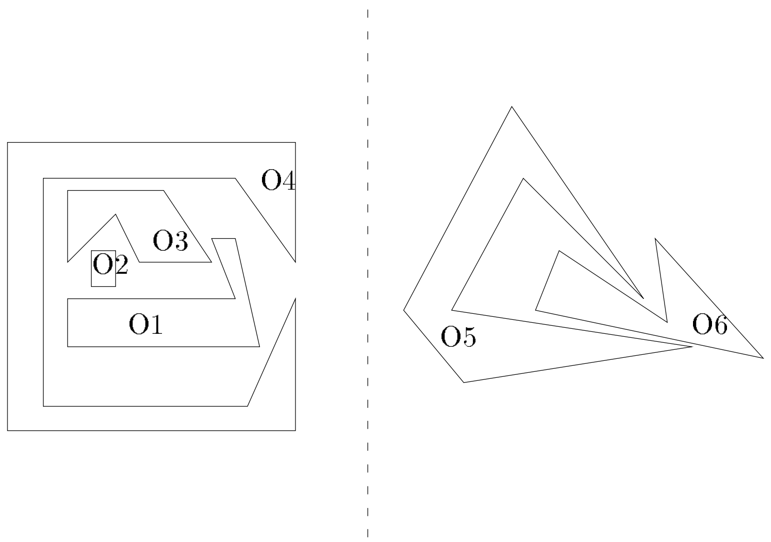

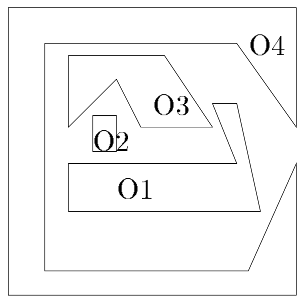

In solving a path-planning problem that only refers to objects o1, o2, o3, and o4, it is possible to ignore objects o5 and o6.

Figure 15: Ignoring separable objects

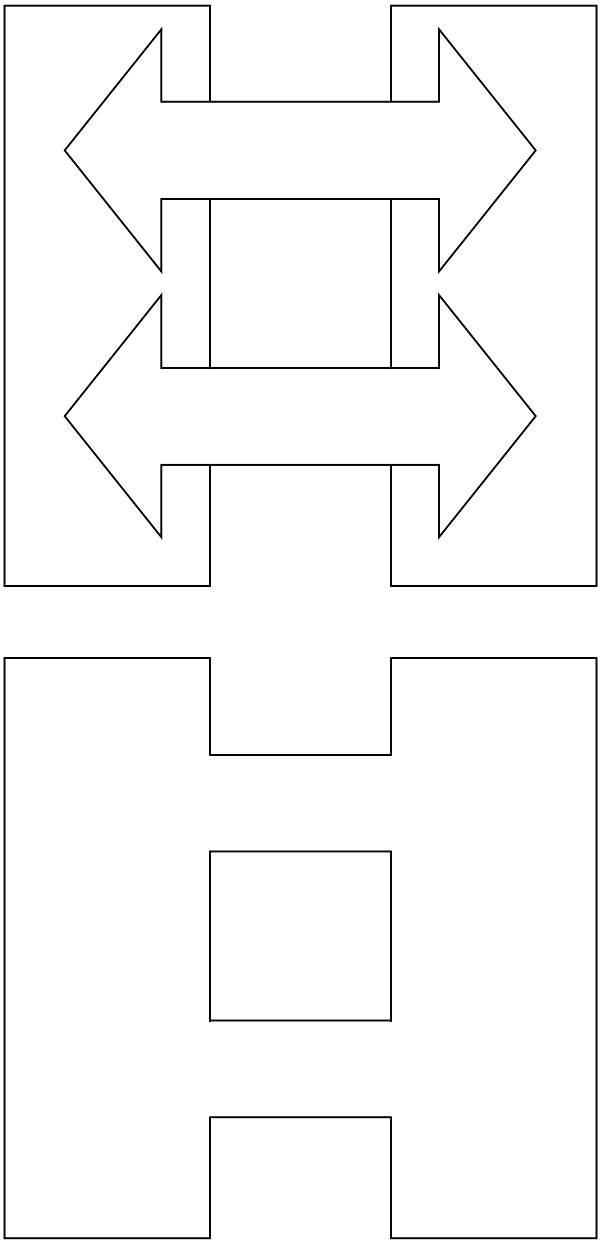

The four tightly bound objects can be treated as a single object.

Figure 16: Coalescing bound objects

and $\mathcal{O}_2$ be the set of objects whose place is a subset of $\mathbf{R}_2$. Let ϵ be the distance between $\mathbf{R}_1$ and $\mathbf{R}_2$; since $\mathbf{R}_1$ and $\mathbf{R}_2$ are compact and disjoint, $\epsilon > 0$. It is immediate that it is possible to translate the objects in $\mathcal{O}_1$ together by a distance less than ϵ and to keep the object in $\mathcal{O}_2$ fixed, without causing any objects to overlap. Thus, if the union of the place is disconnected, then the objects are not tightly bound. ∎

Lemma 4.1: If $\mathcal{O}$ is tightly bound in $\langle D, C\rangle$, then the union over $O \in \mathcal{O}$ of $DC[O]$ is an object-like region.

Proof: Immediate from Lemma 4.1. The other features of an object-like region (closedness, boundedness, regularity, and non-emptiness) are preserved under any finite union.

Theorem 4.3: Let $\langle D, C\rangle$ be a scenario over the set of objects $\mathcal{O}$. Let $\mathcal{O}_B$ be a subset of $\mathcal{O}$ which is tightly bound in $\langle D, C\rangle$. We replace the bound objects $\mathcal{O}_B$ by a single object O_N whose place is the union of the places of the bound objects. (By lemma 4.2, the result is still a consistent scenario.) Formally, let O_N be an object not in $\mathcal{O}$, and let $\mathcal{O}' = \{O_N\} \cup \mathcal{O} - \mathcal{O}_B$. Define the scenario D', C' over $\mathcal{O}'$ as follows:

- For $O \in \mathcal{O} - \mathcal{O}_B$, $C'[O] = C[O]$ and $D'[O] = D[O]$.
- $D'[O_N] = \bigcup_{O \in \mathcal{O}_B} CD[O]$.
  $C'[O_N]$ is the identity mapping.

Then accessible(C, D) is isomorphic to accessible(C', D').

Proof: The isomorphism Λ is defined as follows: Let $C'_1 \in$accessible(C', D'). We define $C_1 = \Lambda(C'_1)$ as follows:

- For $O \notin \mathcal{O}_B$, $C_1[O] = C'_1[O]$.
- For $O \in \mathcal{O}_B$, $C_1[O] = C'_1[O_N] \circ C[O]$

The proof that this is an isomorphism is straightforward. ∎

4.4 Reduction of dimension

If objects are prisms confined to lie in a common plane, then the system can be treated as if it were two-dimensional.

Theorem 4.4: Let $\langle D, C\rangle$ be a scenario in which the place of every mobile object is a prism from $z = 0$ to $z = L$. Suppose, further, that, for any two configurations $C_1, C_2 \in$accessible(C, D) and for any object O, the mappings $C_1[O], C_2[O]$ are related by a motion in the $x - y$ plane. Let C' be the projection of C onto the $x - y$ plane. Then accessible(C') is isomorphic to accessible(C) and is equal to the projection of accessible(C) onto the $x - y$ plane.

Proof: Immediate. ∎

Other reductions of dimension are also possible:

- If the objects are constrained to linear motion, then the system can be treated as if it were one-dimensional.
- If an object is circular or spherical, then the degree(s) of freedom corresponding to its rotation can be ignored.

- If the size of an object O is small as compared to the size of the other objects and to the distances between other objects, then O can be approximated as a point, eliminating the degrees of freedom corresponding to its rotation. If O is a very thin cylinder, then it can be approximated as a line segment, again eliminating the degree of freedom corresponding to rotation about its axis. Point objects and a single linear or laminar object can be accomodated in KRSO if the criterion for overlap is revised to read, "Two objects overlap if any point of one is an interior point of the other." Accommodating multiple linear or laminar objects is more difficult.

5 Approximation of configuration space

A KRSO problem is often addressed by abstracting or approximating the configuration space directly. Indeed, most techniques for reasoning in this domain, including most techniques, exact or approximate, for solving the motion planning problem, do this in one way or another. Our discussion here largely follows [Joskowicz, 89] (a similar study was carried out independently in [Nielsen, 88]) though we include some categories not considered there.

5.1 Partitioning configuration space

A large family of techniques for analyzing KRSO problems, particularly motion planning problems, involves configuration space into regions. The space can then be characterized as an undirected graph, whose nodes are the regions of configuration space, and whose arcs represent bordering relations between regions. This graph is known as a *place vocabulary* [Forbus, 80], [Faltings, 87]. Once the graph is connected, path finding through configuration space can be reduced to path-finding through the graph.

For example, for the system in figure 17 it is natural to divide configuration space into the following nine regions:

A. Four regions where the ball is in contact with one side.

B. Four regions where the ball is in contact with two side.

C. One region where the ball is not in contact with any side.

Figure 17 shows the corresponding graph.

For the system in figure 18 a natural division of configuration space uses 65 regions.

A. 16 regions where the rectangle meets the square at a corner.

B. 16 regions where a side of the rectangle meets the inside of a side of the square.

C. 16 regions where a vertex of the rectangle abuts a side of the square.

D. 16 regions where two vertices of the rectangle abut two sides of the square.

E. 1 region where there is no contact between the rectangle and the square.

Both of the above partitionings are instances of *cellular decomposition,* a general technique for partitioning configuration spaces of objects with semi-algebraic shapes [Mishra, 93]. Other

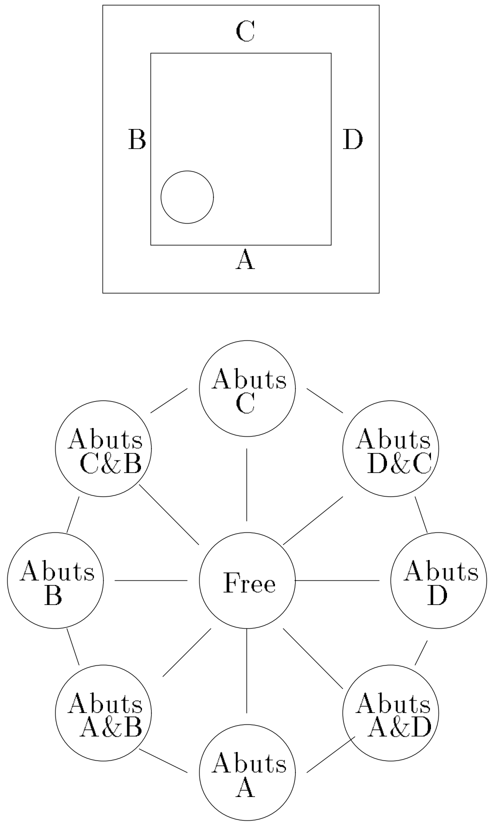

Figure 17: Partioning of configuration space

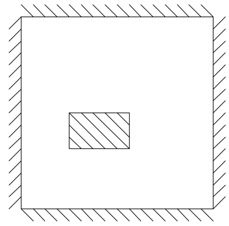

Figure 18: Another configuration space for partitioning

techniques for partitioning configuration spaces for motion planning are discussed in [Ahrikencheikh and Seirig, 94] Of particular interest as regard our discussion here are approximate techniques for partitioning configuration space. [Zhu and Latombe, 91], for example, proposes a generalization of quad-trees where configuration space is hierachically divided into rectangular boxes. A box entirely in free space is labelled FREE; a box entirely in forbidden space is labelled BLOCKED; a box that spans the two spaces is labelled MIXED, and is further decomposed at a more refined level of the hierarchy. The paper presents algorithmic techniques for choosing a partition into boxes that generates many FREE and BLOCKED cells at high levels of the hierarchy.

Once the configuration space has been partitioned, then the graph can be used to answer various motion planning problems. For example, one can infer from the graph in figure 17 that, in order to get from a configuration where the circle borders both A and D to one where it borders both B and C, it is necessary to pass through either a configuration where the circle borders A and B, or one where it borders C and D, or one where it is in the interior.

5.2 Gap closing

Kinematic mechanisms, such as meshed gears, often have a small amount of "play" that can is irrelevant to their global behavior or to their function. In configuration space, these correspond to narrow channels. These can be abstracted to some contained surface (often one of the boundary surfaces of the channel), greatly simplifying subsequent calculations. This is gap-closing at the configuration-space level [Nielsen, 88] [Joskowicz, 89]; it is more generally applicable than gap-closing at the object level, but, of course, can only be applied after the configuration space with its gap has been calculated.

5.3 Periodicity

Configuration spaces of artificial systems are often very regular in structure. For example, if two gears with respectively M and N teeth are meshed, where M and N are relatively prime, then the configuration space consists, intuitively, of MN identical regions, one for each pair of teeth, strung together in cyclic order. It makes sense, then, to describe the geometry of a characteristic region, and state that it is repeated MN times, rather than describe each region separately [Joskowicz, 89]. As far as I know, there has not been a systematic study in the AI literature of representations of this kind of repeated spatial structure and inference using such representations.

5.4 Distinguished points in configuration space

The concrete 15-puzzle is a toy consisting of 15 labelled unit square pieces in a 4×4 frame. The abstract 15-puzzle consists of a 4×4 array, with sixteen different value, one of them "blank", and the constraint that the only possible action is to swap the blank entry with any of its neighbors. (Figure 19)

It is easy to find a mapping of the abstract puzzle into the concrete puzzle. The most natural mapping is to map a value of the array into an "aligned" position of the toy. Swapping the blank with a neighboring entry in the array then corresponds to moving the corresponding piece into the empty space, leaving a blank space behind and forming a new aligned position of the puzzle. Thus it is clear that any sequence of operations on the array can be mirrored in a sequence of movements in the toy.

Note that the abstraction introduces the "blank" as an entity that moves around, where in the

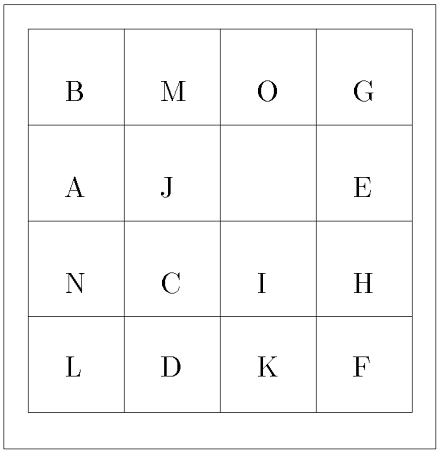

$$\begin{bmatrix} B & M & O & G \\ A & J & \_ & E \\ N & C & I & H \\ L & D & K & F \end{bmatrix}$$

Figure 19: The fifteen puzzle and its abstraction

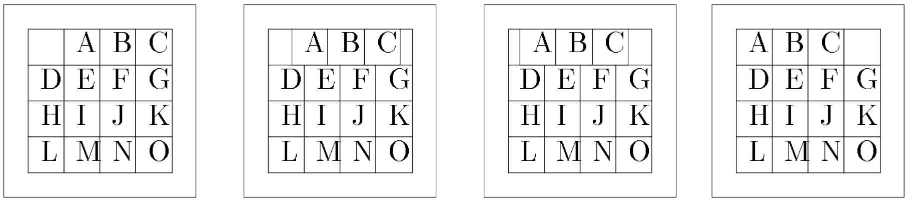

Sliding the top row of the puzzle

_	A	B	C		A	_	B	C		A	B	_	C		A	B	C	_
D	E	F	G		D	E	F	G		D	E	F	G		D	E	F	G
H	I	J	K		H	I	J	K		H	I	J	K		H	I	J	K
L	M	N	O		L	M	N	O		L	M	N	O		L	M	N	O

Corresponding abstraction (see text below)

Figure 20: Sliding squares as a group

geometric description it is just the complement of the places of the pieces. [Casati and Varzi, 94]

The reverse mapping is subtler. The mapping described above is not, of course, surjective; non-aligned positions of the toy do not correspond to array values. Moreover, there are motions from one aligned position to another aligned position that do not correspond to simple swaps of the blank with a neighbor; namely, those where a group of blocks in the row with the blank are slid together. (Figure 20)

Nonetheless, the array problem is a "sufficient" abstraction of the toy by virtue of the following facts:

- The blank space in the toy is always either in a single row or a single column.
- Let ϕ be a feasible history of the toy starting in C_1 and ending in C_2. If C_2 differs from C_1 in either

  a. the placement of the squares outside the incomplete row/column; or

  b. the order of the squares within the incomplete row/column;

  then ϕ contains an aligned configuration.
- Let ϕ be a feasible history of the toy starting in C_1 and ending in C_2, where C_1 and C_2 are aligned configurations and are the only aligned configurations in ϕ. Then the array corresponding to C_1 can be transformed to the array corresponding to C_2 using a maximum of three swaps.

From these three facts it follows that any feasible history ϕ of the toy can be mirrored in a history of swaps in the array where at least every third array value corresponds to an aligned configuration in ϕ. Formally, we have the following:

Theorem 4.4: Let ϕ be a feasible history of the toy over time interval $[0, T]$, where $\phi(0)$ and $\phi(T)$ are aligned. Define the sequence of time points $\langle t_0, t_1 \ldots t_m \rangle$ recursively as follows:

i. $t_0 = 0$.

ii. For $i > 0$, t_i is the least time greater than t_{i-1} such that $\phi(t_i)$ is an aligned configuration different from $\phi(t_{i-1})$.

Then there exist a sequence $\langle A_0, A_1 \ldots A_k \rangle$ of array values and a subsequence $\langle d_0 = 0, d_1 \ldots d_m = k \rangle$ of the indices $0 \ldots k$ such that

a. For $i = 1 \ldots m$, $\phi(t_i)$ is the aligned position of the toy corresponding to A_{d_i}.

b. A_i follows from A_{i-1} by swapping the blank with a neighbor.

c. $d_{i-1} + 1 \leq d_i \leq d_{i-1} + 3$.

It seems plausible to conjecture that this abstraction is found in two stages: first, abstracting from the toy to the sequence of aligned positions $\phi(t_i)$, and observing that successive aligned positions differ by a motion of squares within the incomplete row/column; and, second, performing a transitive reduction, and noticing that any motion of squares in the incomplete row/column is equivalent to at most three motions of one square at a time. Still the abstraction is formally quite complex, and it is remarkable that it is found so automatically by everyone who works with the toy for any length of time.

Alternatively, it is possible to derive this abstraction through a clever partitioning of configuration space (section 5.1) as follows: A configuration of the toy is abstracted as an array where the pieces outside the incomplete row/column have corresponding entries in corresponding places; where pieces in the incomplete row/column have the corresponding entries in the same order; and where the position of the blank entry is determined according to the following rule. (Figure 20)

Position of blank	Minimum coordinate of pieces
Square 1	Piece $1 > 3/4$
Square 2	Piece $1 \leq 3/4$; piece $2 > 3/2$;
Square 3	Piece $2 \leq 3/2$; piece $3 > 5/4$.
Square 4	Piece $3 \leq 5/4$.

Thus, if the pieces are moved in a block from right to left, then the blank is conceptually swapped from left to right. At some stages, the blank will be in places not corresponding to any part of the actual vacancy.

The cellular decomposition of the configuration space (section 5.1) yields a state space with 6-1/2 times as many nodes as the array abstraction. Each cell is a three dimensional manifold. (In general position, the three squares in the row or column with the space are all three separated, so there are three degrees of freedom, one for the motion of each square.) It may then be feasible to automate the solution of the puzzle by searching this partitioning of configuration space, without incurring costs inordinately greater than the usual search techniques for the discrete fifteen puzzle.

Note, however, that the the search space must be built generatively during search, rather than constructed explicitly in full. This is a truism for search of discrete spaces, but rather less common in search of configuration spaces. Such a generative representation is related to, but more complex than, the periodicity abstraction discussed above (section 5.3).

On the other hand, if we slice down the pieces to be .99 unit square, so that the puzzle has a little "give" then all hell breaks loose in configuration space. First, individual cells become forty-five dimensional (three degrees of freedom for each piece). Second, the number of cells increases frighteningly. (I can prove an increase by a factor of 10^{14}, but this is counting very conservatively, and I would estimate that the true value is more than 10^{20}. By way of comparison, the number of attainable array configurations is $16!/2 \approx 10^{13}$.) Running this in the reverse direction is an extreme example of the power of approximation to a tight fit (section 3.1).

5.5 Kinematic device as black box

Mechanisms are often designed as one-degree of freedom systems which generate an output motion from an input motion. For instance a gear train transforms rotational at motion at the input to rotational motion at the output, multiplying angular velocity by a fixed constant. The Peaucellier inversor transforms linear motion to circular motion. If the input and the output are each characterized in terms of a single parameter, then the overall behavior of the system can be characterized by specifying output as a function (possibly multi-valued) of input.

6 Approximation of KRSO by a more complex theory

Somewhat surprisingly, it is sometimes natural or useful to approximate KRSO by theories that are more complex, or at least by theories that are not, in any clear sense, simpler than KRSO.[1] We will consider two such cases:

- The approximation of KRSO by the dynamic theory of motion in a potential field.
- The approximation of a chain by a string.

6.1 Potential fields

The use of potential fields to solve motion planning problems was first studied in [Krogh, 1984] and [Khatib, 1985]). In this technique, the obstacles are viewed as creating a repellent potential field, and the goal is treated as creating an attractive potential field. These fields are then added together, and the motion of the moving objects in the net field is calculated. The advantages of this approximation are that it is often computationally efficient, and that it yields paths that stay away from the obstacles. The chief disadvantage is that it is susceptible to local minima of the potential field, and is therefore incomplete.

6.2 Approximation of a chain by a string

It is sometimes natural to approximate a regular structure of many rigid objects with many degrees of freedom by a single uniform non-rigid object with infinitely many degrees of freedom. In this paper, we will consider a single example: the approximation of a chain as a string. Other examples would include approximating chain mail as cloth, or a pile of sand as a fluid.

The notion of a chain is not very precisely defined; different structures of objects can be more or less chain-like in various respects. For illustration, we will give a quite narrow definition.

[1] I would think that both of the theories described below are in fact more complex, in the sense of being harder to axiomatize in "reasonable" languages, but the point is not a critical one.

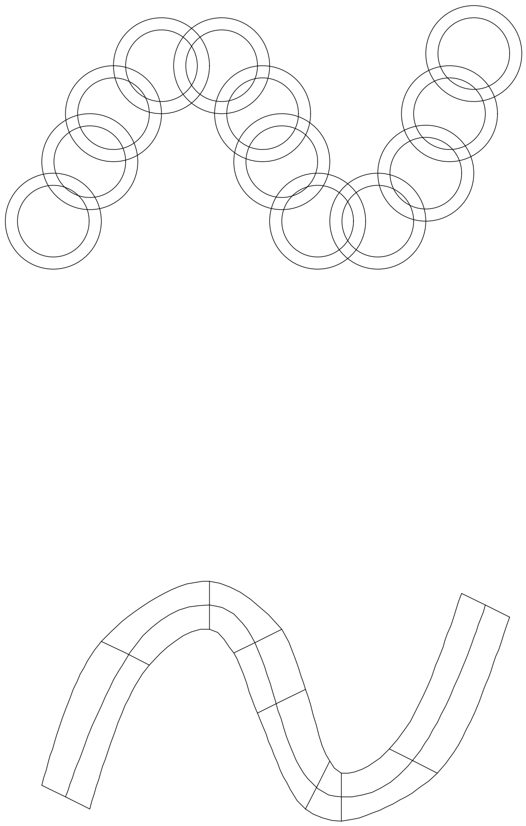

Figure 21: Chain and string

Definition 5.1: Let **C** be a circle centered at the origin in the $x-y$ plane of radius R. Let **T** be the locus of all points whose distance from **C** is no greater than E. Then **T** is a circular torus in standard position. The inner radius of **T** is $R-E$, and the outer radius is $R+E$.

Definition 5.2: Let **T** be a circular torus in standard position whose inner radius I is at least half as large as the outer radius. The *extended chain* with n links of shape **T** is the scenario over a sequence of n objects $\langle O_1 \ldots O_n \rangle$ such that

i. The shape of each object is **T**.

ii. If k is odd, then the placement of O_k is the translation of **T** by the vector $2(k-1)I\hat{x}$.

iii. If k is even, then the placement of O_k is the rotation of **T** by $\pi/2$ around the $\hat{x}$ axis, followed by its translation by the vector $2(k-1)I\hat{x}$.

The *length* of this extended chain is $2(n-1)I$, the distance from the center of the first link to the center of the last link.

Definition 5.3: The *chain* of n links of shape **T** is the zone of configuration space accessible from the fully extended chain of n links. (Figure 21)

Definition 5.4: The *tube* of a chain in a given configuration is the union of the convex hulls of the links.

Definition 5.5: A chain *follows* a curve **C** if

- The curve lies inside the tube of the chain;
- Centers of successive links lie in order on the curve.

Definition 5.6: An *end* of a chain is the point on either the first or last link which is furthest from the convex hull of the next link.

We now turn to the definition of strings. Strings, of course, are non-rigid and thus lie outside KRSO.

Definition 5.7: The *shape* of a string is a right circular cylinder.

Definition 5.8: A *placement* of a string of length L is a twice-differentiable, length-preserving function $\mathrm{f}(X)$ over the interval $[0, L]$. That is, for $X \in [0, L]$, the arc-length of $\mathrm{f}([0, X])$ is equal to X.

This curve is called the *spine* of the string.

Definition 5.9: A string *follows* a curve if the spine of the string is equal to the curve. The *ends* of the string are the ends of the spine.

Definition 5.10: Let $\mathrm{f}(X)$ be a placement of a string of length L and radius R. The *place* of the string under that placement is the set $\{\mathbf{q} + \Delta\vec{N}\}$ where $\mathbf{q}$ is in the image of f; $\vec{N}$ is normal to the curve f at $\mathbf{q}$; and $\Delta \leq R$.

Thus, we place a little disk around each point in the spine of the string, which remains perpendicular to the spine as it moves around. A placement is *feasible* if it does not cause any of these two disks to overlap.

Definition 5.11: A placement $\mathrm{f}(X)$ of a string of radius R is *infeasible* if there exist two distinct points $\mathbf{q}_1$ and $\mathbf{q}_2$; two normals $\vec{N}_1$ and $\vec{N}_2$ to the tangent to $\mathrm{f}(X)$ at $\mathbf{q}_1$ and $\mathbf{q}_2$; and two quantities $\Delta_1, \Delta_2 < R$, such that $\mathbf{q}_1 + \Delta_1\vec{N}_1 = \mathbf{q}_2 + \Delta_2\vec{N}_2$. If a placement is not infeasible, then it is feasible.

Lemma 5.1: If $\mathrm{f}(X)$ is a feasible placement of radius R, then the radius of curvature of f is nowhere less than R.

Proof: This is a basic result from differential geometry.

We can now note that strings and chains have various properties in common:

Theorem 5.1: In the configurations that place the two ends of a string/chain furthest apart, the string/chain follows a straight line.

Theorem 5.2: Let **C** be a closed, smooth simple (non-self-intersecting) curve of finite length L and bounded curvature. Then any sufficiently fine string/chain of length L can be placed so as to follow **C**. ("Sufficiently fine" means that the radius of the string or the outer radius of links is sufficiently small.)

Theorem 5.3: In any placement, a string/chain occupies a constant volume. The volumes of two (strings of the same radius / chains of the same link) are proportional to their lengths.

Theorem 5.4: Let P_1, P_2, P_3 be placements of three (strings of the same radius)/(chains over the same shape of link) such that

1. $P_1 \cup P_2$ and $P_2 \cup P_3$ are feasible placements of a string/chain.

2. P_1 and P_3 are disjoint.

Then $P_1 \cup P_2 \cup P_3$ is a feasible placement of a string/chain.

That is, strings and chains can be strung together. The bridge segment P_2 handles the connection. The theorem can be stated more simply for either strings or chains, but this statement allows them both to be handled together.

The proofs of these and further theorems are given in [Davis, in prep.]

Why, one may ask, should the string be viewed as an abstraction of the chain, rather than *vice versa?* What computational advantage is there to viewing a chain, with finitely many degrees of freedom, as a string, with infinitely many? The answer is that many inferences are easier to perform or to visualize on the uniform structure of the string than on the complex structure of the chain. Consider the following:

Theorem 5.5: Let **C1** be the circle of radius 1 in the $x-y$ plane centered at $\langle 0,1,0 \rangle$ and let **C2** be the circle of radius 1 parallel to the $x-z$ plane centered at $\langle 0,2,0 \rangle$. Let H be a feasible history over two strings/chains O_1 and O_2 such that

1. Throughout H, the beginning and end of O_1 are at $\langle 0,0,0 \rangle$ and the beginning and end of O_2 are at $\langle 3,0,0 \rangle$.
2. At the beginning of H, O_1 follows **C1** and O_2 follows **C2**.

Then at no point during H is there a plane that divides the place of O_1 from the place of O_2.

For a string, this result follows directly from the topological result that if two closed curves move from being linked to being separated, then at some time in between they overlap. (Sketch of proof: The spines are originally linked. If they are later separated by a plane, then in between they intersect at some point, which cannot be an endpoint of both strings. But such an intersection is necessarily infeasible.) The proof for a chain is much more difficult, since there are no particular "spines" of a chain that are barred from overlapping. In fact, there are chain-like structures where the conclusion is false. Consider, for example, two chains, one made of large links where every link has a small gap, and the other made of fine links small enough to fit through the gaps of the large links. Then the fine chain can be passed through the large chain, though the two chains possess many of the other properties of string. It is plausible to conjecture that a system might observe or infer that this property holds for strings, and then reason by analogy non-monotonically that it probably holds for chain-like structures as well. [Davis, in prep.]

7 Approximation of a more complex theory by KRSO

The final category of abstraction we consider is the use of KRSO as an abstraction to more complex theories.

The most important of these is the use of KRSO as an abstraction of the dynamic theory of rigid solid objects (DRSO) [Kilmister and Reeve, 66], [Davis, 88], [Gelsey, 95a]. DRSO extends KRSO by adding such concepts as mass, momentum, force, angular momentum, torque, and energy; positing and constraining normal and frictive forces between objects in contact; and stating constraints on collisions. (Note that DRSO is an entirely different theory from the potential theory considered in 6.1.) KRSO is a strictly weaker theory than DRSO; the ontological sorts, the language, and the theory of KRSO are subsets of the sorts, language, and theory of DRSO. In the terminology of [Giunchiglia and Walsh, 92] KRSO is a theorem-decreasing (TD) abstraction of DRSO; in the terminology of [Nayak and Levy, 94], a model-increasing (MI) abstraction. Thus, if H is a history satisfying DRSO, then the kinematic portion of H satisfies KRSO.

In uncoupled or loosely coupled systems, the behaviors of KRSO are arbitrary paths in a high-dimension configuration space, while the initial-value problem in DRSO usually (though not always) has a unique solution; almost always, the class of solutions is a class of paths of small dimensionality. In cases like these, KRSO offers only weak, though significant, constraints on the dynamic behavior. Typically a purely kinematic analysis of such systems is too weak to be adequate, except in cases where the additional information needed for dynamic analysis (masses, coefficients of friction, coefficients of restitution, exogenous applied forces) is so weak as to add almost no further constraints beyond the kinematic.

KRSO therefore becomes an important approximation in cases where the kinematic constraints are very strong, especially in systems with effectively one degree of freedom. Since many or most mechanisms are built this way [Joskowicz and Sacks, 91], this is a very important special case. In particular, if a given input driving motion has a *unique* kinematic response, then the dynamic system must either respond the same way or jam. I believe that such jamming must be due to friction, though I have not been able to prove it.

Conjecture 6.1: Let Q be a dynamic system of frictionless rigid solid objects. Let H' be a kinematic driving history over a single object O. Suppose that the response space to H' over Q contains a single history H. Then there exists an applied force F on O such that H is the kinematic portion of the dynamic response of Q to O.

The central point of the above conjecture is that, if the system does not jam kinematically, and it is frictionless, then it cannot jam dynamically.

Another conjectured relation between kinematics and dynamics is that any kinematically possible behavior can be achieved dynamically if suitable external forces can applied to all the objects.

Definition 6.1: KRSO+ is the theory consisting of KRSO plus the constraint that every history has a second derivative continuous from the right (i.e. finite acceleration.)

Conjecture 6.2: Let H be a history of a set of objects $\mathcal{O}$ consistent with KRSO+. Then there exists a pattern of input forces whose dynamic response is kinematically equal to H.

If this conjecture is correct, then KRSO+ plus the constraint that accelerations are finite is the best possible kinematic abstraction of DRSO; any behavior consistent with KRSO+ is attainable within DRSO. If it is not correct, then, at least in principle, there are further kinematic constraints; that is, there are some kinematic behaviors consistent with KRSO+ which can never be instantiated in DRSO.

It should be noted that kinematic analysis of a problem is not necessarily simpler than dynamic analysis. An extreme example is prediction of behavior where the starting situation is motionless and stable. In such a case, dynamic analysis requires only determining that the forces are balanced, and predicting that the system remains static, while kinematic analysis may require computing a complex configuration space.

KRSO can also be used as an abstraction for certain problems involving materials that are not rigid, or even materials that are not solid. For instance, let O be a system consisting of a closed bottle half full of liquid. In computing the kinematic interactions of O with external objects, the contents of the bottle may be ignored, using a generalization of theorem 4.1. Another example: Let O be a tightly wound ball of string. Then, as long as O remains wound, it may be treated as a solid object, using a generalization of theorem 4.3.

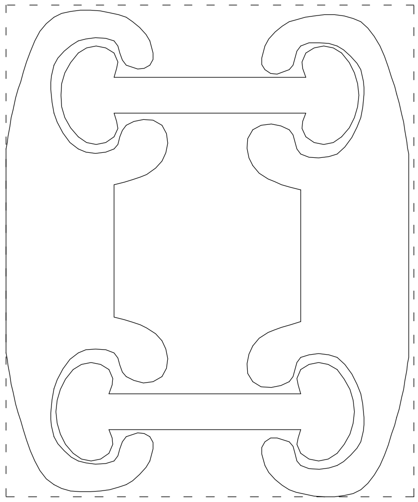

Figure 22: Cascaded approximations

8 Ideas toward implementation

A system using these kinds of abstractions might be build on the following lines: The knowledge base consists of approximations of the problem labelled by some measure of the accuracy of the approximation. (In the case of heuristic rules, this label might be something as vague as "Probably close to correct.") Each abstraction rule is implemented as a production or rewrite rule whose left-hand side is some feature of the approximation and whose right-hand side is a new appproximation with a new measure. Eventually, when the approximate problem is deemed sufficiently tractable, appropriate solution techniques are applied. The accuracy of the answer thus obtained is determined using the accuracy measure on the approximation. In approximation towards simplicity, the starting point would be a precise description of the system; in approximation from ignorance, the starting point would itself be an approximation.

An important feature of such an architecture would be the ability to cascade different approximations together. For example, in figure 22 the system might first approximate the shape of the objects so that they fit tightly, as in section 3.1; then approximate the four tightly coupled objects as a single object (section 4.3); then eliminate the internal cavity (section 3.2); then approximate the shape of the external hull as the dashed rectangle (theorem 3.2). Another example: Starting with a dynamic theory of two meshed gears, first abstract the physics to the kinematic theory (section 7) then approximate the thin configuration space as a one-dimensional configuration space (section 5.2); then abstract the one-dimensional configuration space as a transformer for angular velocity (section 5.5).

The first step in the first series of approximations above illustrates a notable difficulty in the design of such a system: The only function of one approximation (reshaping the objects for a tight fit) may be entirely to achieve the preconditions of another (coalescing tightly bound objects). Such patterns of inference are generally difficult to attain in purely forward-chaining systems. On the other hand, this inference domain does not look at all suitable to any kind of backward chaining.

Designing such a system involves (at least) two major difficulties:

- Many of the individual heuristics are very open ended. For example, using theorem 3.1, one can deduce a free path for a problem from a free path from any set of circumscribed shapes, or deduce the non-existence of a free path from the non-existence of a free path from any set of inscribed shapes. How is a suitable approximation to be chosen?
- The control problem: When more than one heuristic is applicable, how does one choose which to apply? How does one decide when to stop approximating and just calculate on the current approximation?

9 Conclusions: Abstraction in KRSO and elsewhere

Let A and B be two physical descriptions of a situation. For instance, A might describe the shape of object O as the ellipsoid $9X^2 + 4Y^2 + Z^2 + 3XY - YZ - 2XZ - 100 \leq 0$, and B might describe the shape of O as some nearby 600-sided polyhedron. Is it more reasonable for a reasoning program to use A as an approximation for B, or B as an approximation for A, or neither? This depends on

- The problem to be solved. For example, if you are trying to determine whether the point $\langle 1,1,1 \rangle$ is inside the figure, then it will probably be easier to use the ellipsoid than the polyhedron. If you are trying to find the surface area of the figure, then it may well be easier to sum the areas of the faces of the polyhedron than to compute the surface integral over the boundary of the ellipsoid.
- The low-level computing facilities available. Judgments of comparative computational difficulty, such as those in the previous paragraph, depend on the hardware and software facilities available. It makes a difference whether your low-level geometric routines are using Mathematica on a SPARC or an numerical package on a Cray.
- Other aspects of the problem solver itself, including the other approximation techniques it contains and its control structure. As we have seen (section 8), the point of applying approximation technique P may be just that it sets up the preconditions for approximation technique Q. Therefore P will be worthwhile only if the problem solver will also apply Q. This, in turn, depends on whether the problem solver knows about Q and on whether the control regime will allow the problem solver to apply Q in reasonable time.

In the final analysis, then, approximation and abstraction are categories of problem reduction. At this level of generality, it is probably no more reasonable to look for a schema governing all types of approximation and abstraction than it would be to look for a schema governing all types of problem solving.

As far as we have seen in this study, abstraction and approximation in KRSO do not much resemble the elegant metatheoretic structures of [Giunchiglia and Walsh, 92], [Nayak and Levy, 94], or [Weld, 90]. We have seen cases where systems of solid objects are abstracted as strings, and cases where strings are abstracted as solid objects. We have seen heuristics which allow one shape to be replaced by another shape, subject to a wide range of geometric conditions, depending on the characteristics of the problem and the questions being asked.

This is not to suggest that there is no structure to the theory and no guidance for a problem solver, merely that this guidance is partial and local. For example, in purely mathematical terms, theorems such as 3.10, 4.1, and 4.3 run just as well backwards as forwards; we could carve out inaccessible cavities, add inaccessible objects, and split up single objects into systems of tightly bound objects. It is hard, however, to imagine a case where it would be advantageous to do so. In these heuristics, then, there is a clear direction to go from more complex to simpler. But it is

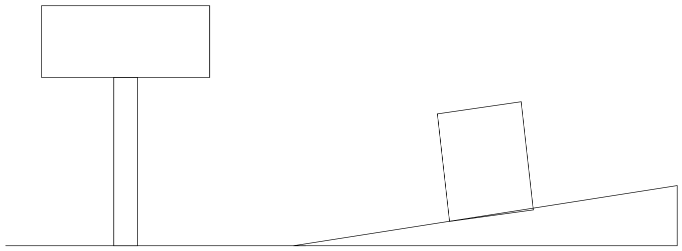

Figure 23: A simple system with friction

not clear that there is any such direction over the space of all approximations and abstractions. In general, it does not seem that you can know in advance whether a given approximation will simplify a problem.

Looking back, this actually applies to some of the theories considered in [Weld, 92] as well. Contrary to the discussion there, there is not necessarily a tradeoff between accuracy and simplicity; in moving from a more accurate to a more idealized theory, we may lose both accuracy and simplicity. Examples:

- Zero friction. If the coefficient of friction in figure 23 is 1/4, then the system is static; if it is zero, then the system will evolve in a complex way.
- Elastic collisions. If the coefficient of restitution is less than 1/2, then the system in figure 24 will, within finite time, settle to an equilibrium state with the ball at the bottom. If collisions are fully elastic, then the long term behavior is more complicated.
- Non-relativistic motion. Consider the following problem: An electron has kinetic energy of 1 Gev. How long does it take to traverse a foot? Relativistically, this does not even require the back of an envelope: The stated kinetic energy is much greater than the rest mass of the electron, so the electron is moving at almost the speed of light = 1 foot per nanosecond. Hence 1 nanosecond is the answer. The non-relativistic answer is not only wrong, but requires more work to find.
- Uniform gravitational field. A projectile is fired out from the earth at 1000 miles per second. What is its state ninety-six hours hence? The correct answer, with an inverse square gravitational law, is that it is very slightly less than 345,600,000 miles away and continuing to move outward at very slightly less than 1000 miles per second. The answer calculated with a uniform gravitational field, ignoring the motion of the earth, is that six hours ago it crashed back into the earth, with consequences that are not easy to predict in detail. If the motion of the earth is considered, then it is somewhere fairly near earth; where, exactly, would take a good bit of work to determine. (The appropriate approximation here, incidentally, is "zero-gravity" but that has to do with the large velocity involved, and would not be easy to find using Weld's techniques.)

(The last two examples above may seem contrived. However, it is not clear what "natural" space of problems is being considered here. In the ordinary course of running a household or designing a mechanism, a physical reasoner would never have to consider either the relativistic correction or the correction for non-uniform gravitational field — they are too small to make any kind of a difference.

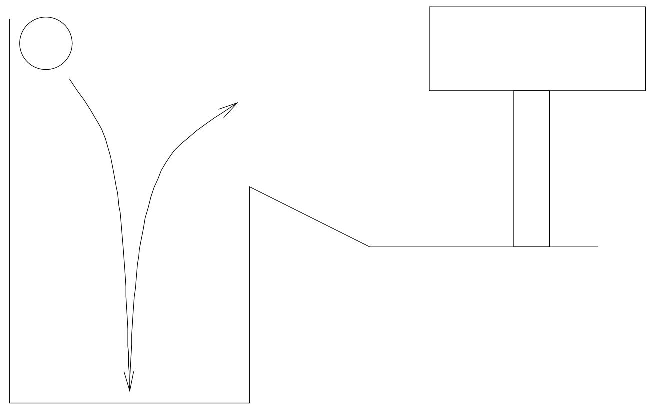

Figure 24: A simple system with inelastic collisions

Therefore, if a reasoner does include such theories as alternates worth considering, the problem space being considered must include some problems with relativistic speeds / heights comparable to the radius of the earth.)

Certainly, there are pairs of theories where the more idealized theory is essentially always simpler to use. One would have to work very hard to contrive problems where it was easier to assume temperature-dependent resistivity than constant resistivity, or where it was easier to use to use the Van der Waals equation than to use the ideal gas law. But this simple relation is not as ubiquitous as Weld's discussion suggests. Certainly, it does not seem to apply to approximation and abstraction in KRSO.

The abstraction framework of [Giunchiglia and Walsh, 92] and the similar proposal of [Nayak and Levy, 94] seem only occasionally relevant here. The one case that fits solidly in these frameworks is the abstraction of the dynamics of solid objects by KRSO. Here the language of the abstraction is a strict subset of the language of the target, and the theorems of the abstraction are a subset (equivalently, the models are a superset) of those of the target. Otherwise, the framework seems less helpful, primarily because these studies are concerned with the abstraction or approximation of one *theory* by another, in the broad sense of theory; while we have found the most common case to be the abstraction of one *problem* by another. The theory often remains the same KRSO, or becomes a more complex theory. Contrary to [Nayak and Levy, 94] it is not the case here that "L_{abs} is a simpler language that makes fewer distinctions than L_{base}." It is the case, certainly, that the sound use of approximation, like any sound inference, is, by definition, a theorem-decreasing (model-increasing) or theorem-constant operation. If one replaces "Object O has shape $\mathbf{Q}$" by "The Hausdorff distance from the shape of O to $\mathbf{Q}'$ is less than ϵ," then the class of models increases, as long as $\mathbf{Q}$ does, in fact, satisfy the given constraint.

To conclude: In KRSO, there are many different techniques for abstraction and approximation that can be used as tools in reasoning. Some of these approximate geometric detail, some simplify the structure of the problem, some map the problem to simpler theories, some map the problem to simpler problems in more complex theories. A library of such techniques could potentially greatly increase the power of a reasoner. We are encouraged, therefore, to explore the power of such techniques and the design of such a reasoner further. We do not, however, see any overarching metalogical structure that sheds useful light on the relation between these approximation techniques.

We suspect that the same will be found to hold in many domains of physical reasoning.

References

Sanjanay Addanki, Roberto Cremonini, and J. Scott Penberthy, "Reasoning about Assumptions in Graphs of Models," *Proc. IJCAI-89*, pp. 1432-1438.

Cherif Ahrikencheikh and Ali Seirig, *Optimized-Motion Planning,* John Wiley, 1994.

Roberto Casati and Achile Varzi, *Holes,* MIT Press, 1994.

Ernest Davis, "A Logical Framework for Commonsense Predictions of Solid Object Behavior," *Int. Journal of AI in Engineering,*, vol. 3, no. 3, pp. 125-140.

Ernest Davis, "Approximations of Shape and Configuration Space," NYU Computer Science Tech. Report #703, 1995.

Ernest Davis, "Strings and Chains: Reasoning about Ensembles using Deduction and Analogy." In preparation.

Brian Falkenhainer, "Ideal physical systems," *Proc. AAAI-93*, pp. 600-605.

Boi Faltings, "Qualitative Kinematics in Mechanisms," *Proc. IJCAI-87*, pp. 436-443.

Rudolf Fleischer, Kurt Melhorn, Gunter Rote, Emo Welzl, and Chee Yap, "Simultaneous Inner and Outer Approximation of Shape," *Algorithmica,* vol. 8, 1992, pp. 365-389.

Kenneth Forbus, "Spatial and Qualitative Aspects of Reasoning about Motion," *Proc. AAAI-80,* pp. 170-173.

Andrew Gelsey, "Automated Reasoning about Machines," *Artificial Intelligence*, vol. 74, 1995, pp. 1-53.

Andrew Gelsey, "Intelligent Automated Quality Control for Computational Simulation", *Artificial Intelligence for Engineering Design, Analysis, and Manufacturing,* 1995, to appear.

Fausto Giunchiglia and Toby Walsh, "A theory of abstraction," *Artificial Intelligence*, vol. 57, pp. 323-389, 1992.

Leo Joskowicz, "Shape and Function in Mechanical Devices," *IJCAI-87,* pp. 611-615,

Leo Joskowicz and Sanjaya Addanki, "From Kinematics to Shape: An Approach to Innovative Design," *Proc. AAAI-88.*

Leo Joskowicz, "Simplification and Abstraction of Kinematic Behaviors," *Proc. IJCAI-89*, pp. 1337-1342.

Leo Joskowicz and Elisha Sacks, "Computational Kinematics," *Artificial Intelligence,* vol. 51, 1991, pp. 381-416.

Leo Joskowicz, Elisha Sacks, and Vijay Srinivasan, "Functional Kinematic Tolerancing," *Computer-Aided Design,* 1996, to appear.

Leo Joskowicz and Russell Taylor, "Interference-Free Insertion of a Solid Body into a Cavity: Algorithms and a Medical Application," *International Journal of Robotics Research,* MIT Press, 1994.

O. Khatib, "Real Time Obstacle Avoidance for Manipulators and Mobile Robots," *1985 IEEE Conf. on Robotics and Automation*, 1985.

C.W. Kilmister and J.E. Reeve, *Rational Mechanics,* American Elsevier, 1966.

B. H. Krogh, "A Generalized Potential Field Approach to Obstacle Avoidance," in *Robotics Research: The Next Five Years and Beyond,* 1984

Matthew Mason, "On the Scope of Quasi-Static Pushing", *Proc. 1985 Third Int. Symp. on Robotics Research,* MIT Press, 1985.

Bhubaneswar Mishra, *Algorithmic Algebra,* In Texts and Monographs in Computer Science Series, Springer-Verlag, New York, 1993. XIV, 416 pp., 9 illus.

P. Pandurang Nayak, "Causal Approximations" *Artificial Intelligence* vol. 70, 1994, pp. 277-334.

P. Pandurang Nayak and Alon Levy: "A Semantic Theory of Abstractions: Preliminary Report," *Workshop on Theory Reformulation and Abstraction,* 1994.

Paul Nielsen, "A Qualitative Approach to Mechanical Constraint," *Proc. AAAI-88*, pp. 270-274.

Raman Rajagopalan, "A Model for Integrated Qualitative Spatial and Dynamic Reasoning about Physical Sytems," *Proc. AAAI-94,* pp. 1411-1417

Aristides A.G. Requicha, "Towards a Theory of Geometric Tolerancing," *The International Journal of Robotics Research,* vol 2, no. 4, 1983, pp. 45-60.

J.T. Schwartz, M. Sharir, and J. E. Hopcroft (eds) *Planning, Geometry, and Compexity of Robot Motion,* Ablex Pubs, 1987.

Daniel Weld, "Approximation Reformulations", *Proc. AAAI-90*, 1990, pp. 407-412.

Daniel Weld, "Reasoning about Model Accuracy," *Artificial Intelligence,* vol. 56, 1992, pp. 255-300.

David Zhu and Jean-Claude Latombe, "New Heuristic Algorithms for Efficient Hierarchical Path Planning," IEEE Trans. on Robotics and Automation, vol. 7 no. 1, 1991, pp. 9-20.