Physics Problem Solving Using Multiple Views

Gordon S. Novak Jr.
Agustin A. Araya

TR-173

Computer Science Department
University of Texas at Austin
Austin, Texas 78712

March, 1981

This research was supported by NSF Award No. SED-7912803 in the Joint National Institute of Education - National Science Foundation Program of Research on Cognitive Processes and the Structure of Knowledge in Science and Mathematics.

ABSTRACT

The primary task in solving a physics problem is to select ways of viewing the problem in terms of physical systems whose behavior is described by physical laws. The physical systems are, in general, only approximate models of the real-world systems. As models of real-world systems are made more accurate, the equations involved quickly become unmanageable; furthermore, certain special cases of physics problems which frequently occur can be solved using greatly simplified equations. Success in solving physics problems therefore depends crucially on selection of physical systems which satisfactorily model the real-world systems and which possess tractable mathematical models. Selection of the physical systems used as models is generally based on qualitative features (including ranges of numerical magnitudes) of the real-world systems. Many problems are best solved by considering the same real-world system from multiple viewpoints and relating the viewpoints to each other, often by identification of their components.

1.Ø INTRODUCTION

Every real-world problem presents the problem solver with potentially unlimited complexity. Given a problem (say, "A car is acted upon by a force; what is its acceleration?"), it is always possible to identify some factors which have been left out of any given analysis of the problem. For example, some energy is consumed by rotational kinetic energy of rotating parts of the car; there are many sources of friction; some of the force may be expended in deforming the car rather than accelerating it; and so forth. The problem solver cannot possibly take all the applicable factors into account. At the same time, it is not possible to determine a priori which factors will have to be considered in a problem of this type; that will depend on the goals of the analysis, the accuracy required, and the magnitudes of certain quantities. The student in an elementary physics course will surely be justified in ignoring all of the complicating factors and simply applying Newton's law, f = ma. On the other hand, an automotive engineer may wish to insure that for forces of certain magnitudes the deformation of the car will absorb enough energy to protect the passengers from injurious acceleration in the event of a collision.

The nature of a physics problem is determined not only by the statement of the problem itself, but also by the implied physical context in which the stated problem occurs. In some cases, a single problem statement can be made into radically different problems (in terms of the physical principles considered in analyzing each problem) simply by changing the numbers in the problem statement. Consider the following ([1], p. 67):

> A rifle with a muzzle velocity of 1500 ft/s shoots a bullet at a target 150 ft away. How high above the target must the rifle be aimed so that the bullet will hit the target?

This problem, as stated, can be solved mentally in a few seconds by an expert in physics. The expert assumes that the motion of the bullet will be nearly linear, calculates the time of its travel (Ø.1 second; clearly, the two numbers were intended to cancel easily), and finally calculates the vertical distance the bullet will fall during that time (1/2*g*t**2, or Ø.16 ft); aiming that far above the target will cancel the fall. This solution is very accurate for this problem (the error compared to the next more complicated model is about one part per million). However, this solution depends on assumptions about the context of the problem -- assumptions which are reasonable for this problem, but not for the whole class of similar problems. These assumptions include absence of air friction, nearly linear motion (so that the time of travel can be calculated independently of the amount of fall), a flat earth, and uniform gravity. Some of these assumptions can be invalidated merely by changing the numbers in the problem; Figures 1-3 illustrate the resulting sequence of problems. If we put the target farther away, the assumption of nearly linear motion will become untenable, and the "rifle bullet problem" will be converted to a "cannonball problem"; however, the assumptions of flat earth and uniform gravitation may be retained. The "cannonball problem" is harder to solve, requiring about half a page of trigonometric equations. If we put the target still farther away, and give the bullet sufficient speed to reach the target, the "cannonball problem" will be converted to an "ICBM problem". Now, we must take into account the curvature of the earth and the fact that gravity will change in both magnitude and direction as the bullet travels along an elliptical path. This more complex solution still ignores air friction, Coriolis forces due to the rotation of the earth, gravitation from the moon, and so forth. If these factors are to be taken into account, it will not be possible to solve the problem in closed form; instead, it will be necessary to numerically integrate a complex set of differential equations.

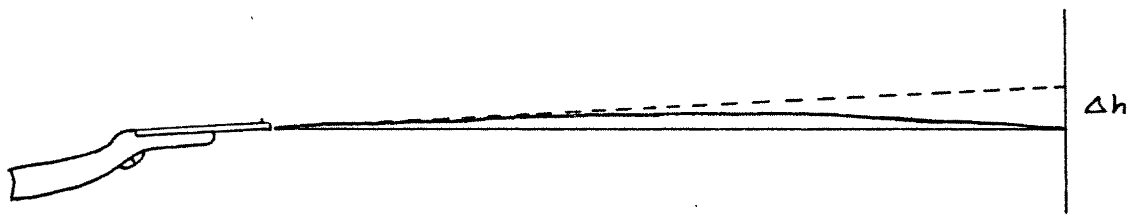

Figure 1: Rifle Bullet Problem (Nearly Linear Motion)

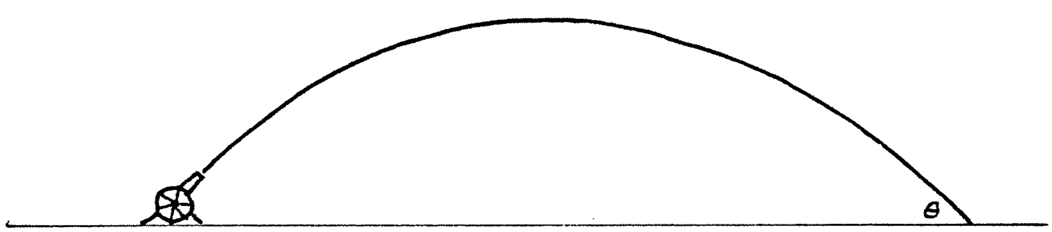

Figure 2: Cannonball Problem (Nearly Parabolic Motion)

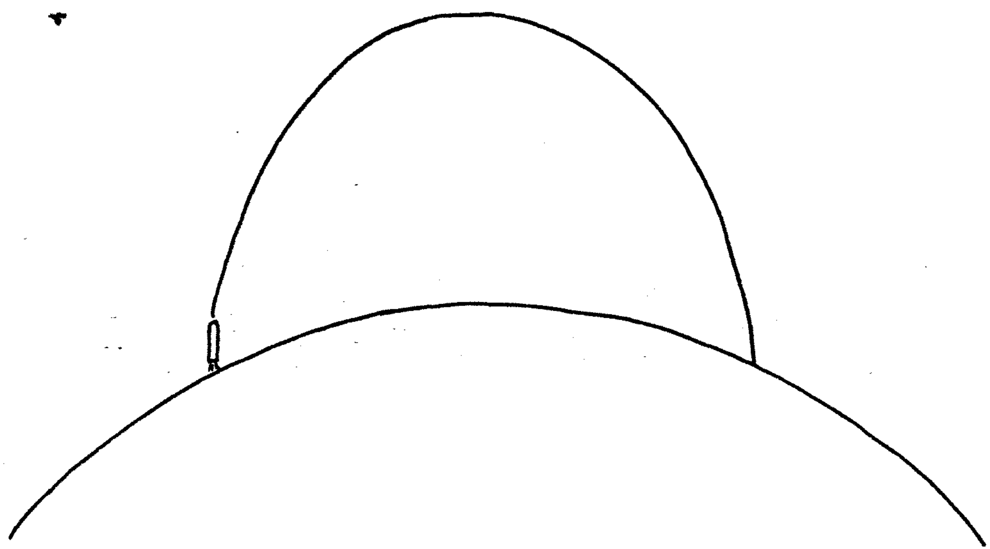

Figure 3: ICBM Problem (Nearly Elliptical Motion)

2.0 EXPERTISE IN SOLVING PHYSICS PROBLEMS

As discussed above, every real-world problem presents the problem solver with potentially unlimited complexity. However, the ability to analyze such problems with sufficient accuracy in real time with limited computational resources is essential to the survival of intelligent organisms in a hostile environment (e.g., students taking an hour exam). The problem solver who knows many physical principles cannot adopt a deductive strategy of applying all the physical laws that are applicable to a given problem. As we have seen, vastly many laws are applicable to every problem; moreover, these laws demand data which are difficult or impossible to obtain. [To analyze Coriolis forces on a bullet, it would be necessary to know the latitude of the gun and the direction to the target.] At the same time, the problem solver must be able to apply all the laws which significantly affect the solution of the problem, or the answer will be wrong. [Failure to consider all the applicable laws is a frequent cause of error in both novice and expert problem solving.] Finally, since there are multiple ways of solving a given problem, some of which are much easier than others, the problem solver must select the easiest method of analysis which will adequately solve the given problem. [The "rifle bullet problem" can be solved using the "ICBM problem" method, but the latter is orders of magnitude more difficult and produces essentially the same answer.]

Solving a physics problem is significantly different from proving a mathematical theorem. In mathematics, knowing only a few facts about mathematical entities (say, that X and Y are integers and $X > Y$) allows one to make true deductions about those entities independent of the context in which they occur (e.g., $X+1 > Y+1$). In physics, however, knowing a few facts (e.g., that a body of mass m is acted on by a force f) does not permit us to make deductions from those facts independent of context (e.g., to deduce that the acceleration of the body is given by the equation $f = ma$). The reason for this is that Newton's law, which we write so simply as $f = ma$, actually relates the acceleration of the body relative to an inertial reference frame to the vector sum of all the forces on the body. In order to find the acceleration of the body using the law $f = ma$, we must verify that there are no other forces on the body (which is unprovable and, in any actual instance, untrue) or assume that the other forces are insignificant; likewise, we must verify or assume that the reference frame is inertial. The laws of physics, being universal, bind the whole universe together and relate every object to every other object. In order to solve any physics problem, we must treat physical systems as nearly decomposable systems [2], ignoring all but a few of these connections. Any application of a physical law thus carries with it a set of implicit ceteris paribus assumptions; whether these assumptions are justified by careful thought and experiment or overlooked in ignorance, they are nonetheless part of the physical law.* Thus, it is clear that it is not possible to solve physics problems by deductive methods in which the laws of physics are stated as theorems unless the system is restricted to operation in a microworld in which the ceteris paribus assumptions of all the laws the system uses are simultaneously satisfied.**

What the problem solver needs is Expertise, which we will define as "the ability to design and execute an effective method of analysis for a given problem." We consider an effective method to be one which will adequately solve the problem at a reasonable computational cost. [Reasonable cost is not merely desirable, but a central issue; intelligence, like politics, is the art of the possible.] The expert will often have redundant methods for solving the same kinds of problems with varying degrees of accuracy and at different costs; thus, the method to be chosen for solving a particular problem will depend partly on the accuracy required of the solution. The central point of our definition of expertise is that problem solving is fundamentally a design problem. The equations which appear in the solution of a physics problem are consequences of the design of a method of analysis for that problem, rather than the primary components of the analysis. Novices (and many experts) tend to give the equations primary status, and novices often solve problems using an "equation-driven search", working backward from the desired unknown through equations which relate it to known quantities

----

* The history of physics provides many examples of such assumptions. Newton's laws, for example, were thought for many years to be unqualified truths before their implicit <u>ceteris paribus</u> assumptions (e.g., speeds much less than the speed of light) were discovered. We still teach and use Newton's laws, but now state the assumptions explicitly.

** It is possible to write a deductive system which will emulate the problem solving methods we are proposing. However, rather than having theorems which state laws of physics as universal truths, such a system would have theorems stating how the problem solver ought to approach problems of particular types. The semantics of such a logical system would specify what "mental states" of the problem solver were possible (rather than what states of the universe are possible), and would not guarantee that the problem solver would not make mistakes (i.e., derive answers which do not conform to physical reality) or derive different answers by approaching a given problem in different ways.

[3]. However, while this method can solve easy textbook problems, it cannot solve many real-world problems and more difficult textbook problems. As in other areas of design, there exist standard designs to be used for certain frequently occurring problems; the easy textbook problems serve to teach standard designs along with examples of problems for which these designs are appropriate. More complex problems, however, require creative design of new methods of analysis.

Larkin, McDermott, Simon, and Simon [3] have reviewed differences between expert and novice humans in solving problems. While novices tend to work backwards using an equation-driven search, experts tend to work forwards from given data (sometimes calculating intermediate results) to a solution, with little or no backtracking. Larkin, McDermott, Simon, and Simon propose that an expert is able to recognize a large number of patterns (on the order of the 50,000 chess patterns which a chess master is believed to recognize) which guide the interpretation and solution of a problem. We share this view; in the sections below, we discuss its implications for an expert program for solving physics problems, and contrast this view with the deductive approach to problem solving.

2.1 Representation

There is no single, "fundamental" level of representation in physics; physics problems range from the level of quarks to particles to atoms to molecules to objects to planets to suns to galaxies to the universe as a whole. For some problems, it is necessary to collect several objects into a single system which is considered as a whole (e.g., the Jugglo problem [4] is best solved by considering a juggler and the balls he is juggling as a single system). In other cases, it is necessary to decompose a single object into components; sometimes the decomposition needed is a "gedanken" one which is chosen to fit a particular problem, rather than being a "natural" decomposition of the object (as in the chain problem to be discussed later).

Multiple representations of the same kind of object are often needed; indeed, many physics books begin with a chapter on vectors and decomposition of vectors into components. Specialized representations are important because certain representations possess invariants which greatly simplify analysis.* Selection of representations and conversion between representations is thus an important component of problem solving; it is something which experts do fluently but novices do poorly [5].

Many problems require that a single object be viewed in

----

* Imagine working a problem involving a swinging pendulum in spherical coordinates with arbitrary zero point and arbitrary orientation of the axes of the coordinate system. As the pendulum swung, each of the three coordinates would change according to a complex nonlinear function. However, if we choose two-dimensional polar coordinates in the plane of the pendulum and with a zero point at the pivot of the pendulum, only a single coordinate (the angle of the pendulum) changes.

multiple ways. Consider the following problem:

> What force applied to the brake pedal is required to stop a 5000 lb car going 55 mph in 5 seconds?

This problem requires two views of the car. In the first view, we trivialize the car and view it as a point mass, and compute the force needed to stop a 5000 lb mass going 55 mph in 5 seconds. In the second view, we expand the car to view its braking system, a complex collection of levers, hydraulic systems, and friction brakes. Sussman and Steele [5] have found that multiple views of electronic circuits are essential to analysis of complex circuitry and have implemented a program for circuit analysis using multiple views (which they call SLICES).

2.2 Reasoning

The expert problem solver possesses a rich set of specialist methods for solving particular types of problems which frequently occur. These methods are often redundant (e.g., the sequence of rifle problem, cannonball problem, and ICBM problem methods), and they may be mutually "contradictory" in the sense that they give different answers to the same problem. In some cases the differences will be small, but in others they will be significant, and some methods will be judged clearly "wrong" for such problems.

The expert's methods must be considered Knowledge rather than universal truths in the sense of physical laws or mathematical theorems. First, the expert's methods operate at a particular level of abstraction, and often are not deductively related to other levels. For example, the law which states that a metal bar expands linearly with increasing temperature was formulated and is used without reference to the underlying atomic level at which the phenomenon can be further explained. Second, the expert's methods are approximate and limited in scope. Thermal expansion is not really linear, but nearly linear over a range of temperature.* Finally, an expert's method always involves a separation of a phenomenon from most of its context, and thus always carries with it ceteris paribus assumptions that such a separation is legitimate.

If a problem solver has a large number of redundant methods, how can the appropriate methods (and only those) be invoked? Clearly, it is inappropriate to use the most complex analysis method to verify that the answer will be in a range for which an easier method could have been used; error measures, which are in fact the precise solution minus the approximate solution, won't work either. Thus, the method to be used must be chosen on

----

* Experts often do not know the boundaries of applicability of their methods. They will use appropriate methods for the problems with which they are familiar; however, if given an unusual problem, they may use inappropriate methods without realizing it ([6], p. 208).

qualitative grounds, which will include ranges of numerical values. (Scientific terms sometimes specify such ranges, e.g. "sub-micron particle".) In some cases, applicability of methods may be explicitly ruled out (e.g., "don't try to differentiate a noisy time function"). Rather than having a small number of universal laws, the expert will have a large number of specialist methods, each applicable in a limited area and governed by a large set of preconditions for applicability. In this sense, expert knowledge is "additive": a new method can be added to the existing set without causing much trouble, because its large set of preconditions keeps it from interacting with much existing knowledge.* A large set of problem solving methods based on Newtonian mechanics can survive the addition of relativistic physics simply by enclosing the whole set of Newtonian methods in a set of extra preconditions. Since there will be redundant methods for solving special cases more easily than more general methods do, and since both the specialist method and the general method will be applicable to the same problem, there must be a mechanism for choosing one (and only one) of the methods. Choosing a problem solving method is similar to parsing sentences in the sense that local ambiguity must be resolved in such a way that the local interpretations fit together into a global interpretation.

----

* In mathematical theorem proving, knowledge is not additive: a single axiom, such as Euclid's fifth postulate, can have far-reaching consequences.

In many cases, the choice of a problem solving method is implied by the choice of views to be used for objects. For example, when we choose to view an accelerating car as a point mass, we are implicitly saying that other effects (e.g., deformation of the car) will be ignored and enabling methods which are specialized to deal with point masses. Making a view thus constitutes a decision that for the given context and analysis goals, an object can be viewed in a certain way, which will have the effect of decomposing the object from much of its environment. It is at this decision point that the _ceteris paribus_ conditions are tested; once the view has been made, these conditions will be assumed implicitly to hold. If in fact some of the preconditions for the selected view turn out to be violated after the view has been created, even the expert analyst may not notice the violation, and the resulting analysis may be fallacious. (This phenomenon is a serious flaw in existing methods of nuclear reactor safety analysis.)

In summary, the approach we are proposing can be contrasted with the way deductive problem solvers have typically been written as follows. Deductive systems have employed deep chains of reasoning from relatively few principles (physical laws stated as theorems); the principles consider a relatively narrow context (have few preconditions); search is used as the method of bringing together related principles (through unification). Our approach emphasizes relatively shallow chains of reasoning, employing a large number of specialist principles; the principles conceptually consider a large context (have many preconditions); search for contradiction in the sense of resolution is not possible or meaningful; the principles to be applied are selected by recognition of special cases.

3.Ø WHAT'S IN A PHYSICS PROBLEM?

Textbook physics problems vary tremendously in difficulty.

> A body of mass 2 kg is acted on by a force of 6 newtons. Find the acceleration.

This problem ([8], p. 39), appearing at the end of the chapter in which Newton's law f = ma is introduced, can be solved entirely at a syntactic level with little understanding of underlying principles or of the assumptions involved (e.g., lack of other forces such as friction). A slightly harder problem type involves two equations rather than a single one (e.g., the above problem with the question "How far does it move in 5 seconds?"); these can be solved by an equation-driven search, still "close" to the syntactic level.

At the opposite extreme are problems like ([9], p. 85):

> If the polar ice caps were to melt, what would happen to the earth's period of rotation?

This problem is difficult in several respects. It is minimally specified in terms of real-world objects which are not "close" to the physical system models required; a great deal of world knowledge must be brought in by the problem solver; the problem is not nicely bounded in terms of the phenomena which may have to be considered. The main difficulty with such a problem is deciding how to "set it up", i.e., how to design the analysis method. Clearly, no "search" method could hope to find a deductive path through known equations to solve such a problem.

These problems illustrate the wide variation in difficulty of physics problems. A number of textbook problems can be solved (by humans or programs) at the syntactic level, or at the syntactic level with a little search, particularly if the problem domain is narrow. Of course, it is easiest to write programs to solve such problems, and such problems have served as a legitimate starting point for problem solving research; however, solving such problems still leaves us far from the goal of solving "real" physics problems, i.e., problems of the types faced by practicing engineers and physicists. The real challenge is to write a program which can solve all the problems in [1] or [9]; this is a difficult challenge indeed. Some of the features which make problems difficult, and which should be attacked by problem solving programs, are listed below.

1. English and pictorial input. Many problem solvers have started with the "simulated output of a parser" or "simulated diagram". However, it is difficult to design a "neutral" intermediate language; care is needed to insure that the "simulated parser" does not inadvertently contain a "simulated physicist".

2. Significant geometry. Most existing programs have used trivial geometry. However, selection and manipulation of geometric models is a central part of physics problem solving. Physics texts are full of diagrams; engineering students score significantly higher than average on tests of visual/spatial ability [17].

3. Many possible views of objects. If an object always plays the same role in problems, the problem of choosing the proper view(s) for that object is trivialized; choosing proper views is the central problem.

4. Minimally specified problems which require assumptions and use of world knowledge for their solution.

4.0 PROBLEM SOLVING USING MULTIPLE VIEWS

The notion of modeling a real-world system by a physical system whose behavior is governed by physical laws is related to the notion of isomorphism in mathematics; Hofstadter [10] mentions this relation frequently. However, our notion of models is looser than the notion of isomorphism (because the modeling is only approximate) and more complex (because different kinds of objects and more complex relationships are modeled). Figure 4 illustrates a simple model in which the combination of two oranges and two oranges in the real world is modeled by the addition of 2 and 2 in the analysis model.

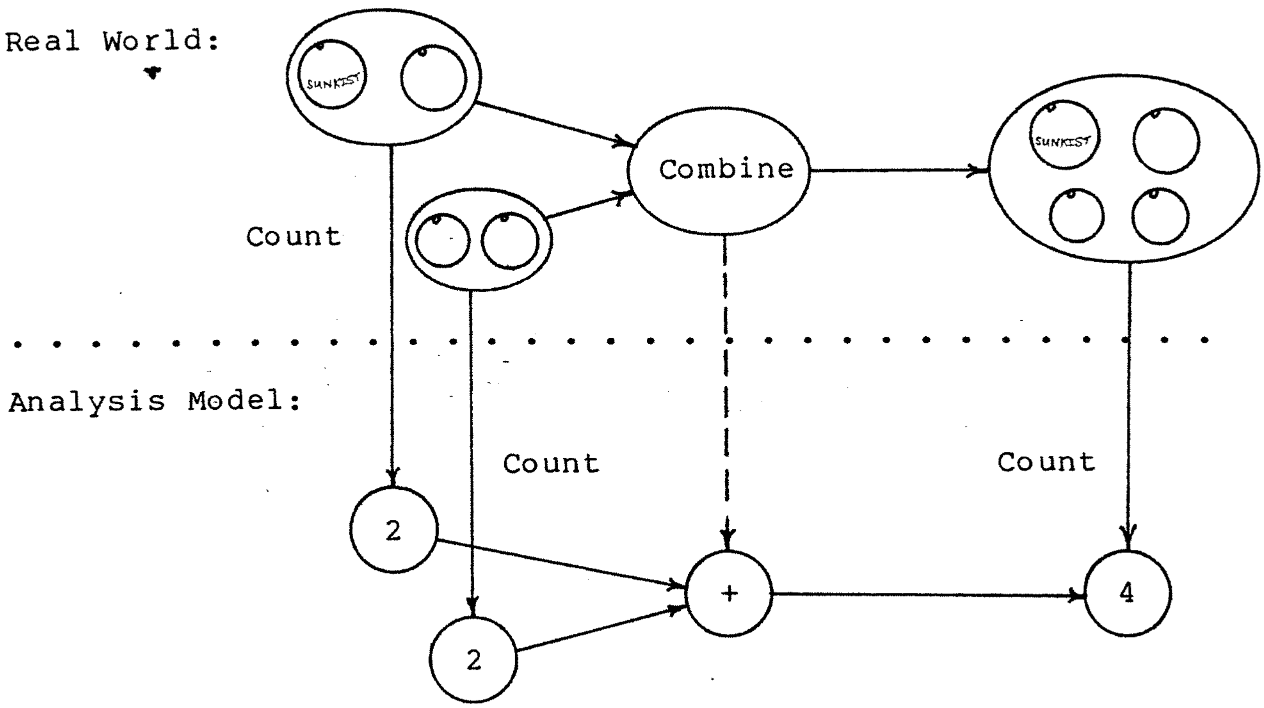

Figure 4: Two Oranges and Two Oranges Makes Four Oranges

This diagram looks very much like the diagrams given for "isomorphism" in mathematics texts. (The bottom layer looks much like the networks of CONSTRAINTS [5] and THINGLAB [11].) The isomorphism consists of two things: a mapping between objects in one set and those in the other set (in this case, counting of oranges) and a correspondence between the operations (combination corresponds to addition). Our orange model, however, is really more complex than an isomorphism. One set of objects is non-mathematical (the oranges); use of this model presumes that there is a way of counting oranges which is satisfactory for our purposes (e.g., how would we treat rotten oranges, tiny oranges, plastic oranges, etc.?). There are also restrictions on the operation: the combination must not smash the oranges, or take so long that they rot, etc. Even this very simple model has its *ceteris paribus* conditions. Thus, we see that the count of oranges is not an intrinsic property of the oranges, but a view of them* [other views are possible, e.g., weight, volume in bushels, etc.]; likewise, the addition operator is a view of the combination process in terms of the views of the oranges.

An isomorphism is specified by two items, the mapping between objects and the operator correspondence. In our models, we will have to keep additional information about the correspondences. A physical law typically relates different kinds of objects using a different view for each; there are often mutual restrictions among the objects ("you can't add apples and oranges"), and other restrictions which apply to the model as a whole. Thus, to use Newton's law f = ma to model a car which is being accelerated by a rocket engine, we view the rocket engine as a force, view the rocket engine and car as a composite object which we then view as a point mass, view the car's environment as an inertial reference frame, view the car as a location relative to the reference frame, and relate the second derivative of the location to the force and point mass using the equation f = ma. Most existing problem solvers (including ISAAC [12,13,14]) have been much too quick to say an object or relation *is* its view and replace it with a variable or equation which has lost its roots; this limits the extensibility of such systems when, for example, we wish to examine the behavior of the helium balloon inside the accelerating car. Failure to remember the basis of an analysis model also makes it prone to error. The constraint networks of [5] and [11] are fine pieces of work, but the networks do not represent their limits of applicability. If, for example, a constraint network were set up to represent the "rifle bullet problem" analysis, and the numbers were changed to those of the "cannonball problem", both of these constraint network systems would blissfully calculate wrong answers. An intelligent constraint network should continually check itself and take corrective action if its bounds of applicability are exceeded. Human experts often check intermediate results for "reasonableness" during problem solving [6].

The use of multiple views in problem solving is nicely illustrated by the following problem ([1], p. 295):

> A flexible chain of weight W hangs between two fixed points, A and B, at the same level. Find (a) the vector force exerted by the chain on each end point and (b) the tension in the chain at the lowest point.

This problem appears at the end of the chapter "Equilibrium of Rigid Bodies"; nowhere does the chapter consider flexible objects. No equations involving chains have been given. However, the problem can be solved using multiple views of the chain as a rigid body; this encourages the student to broaden her definition of rigid body (i.e., a body of invariant geometry). The resulting analysis is illustrated in Figure 5.

Given Problem:

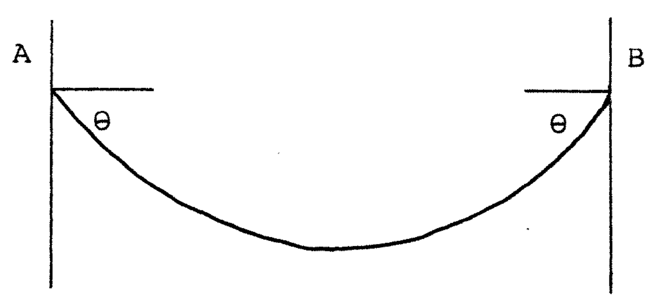

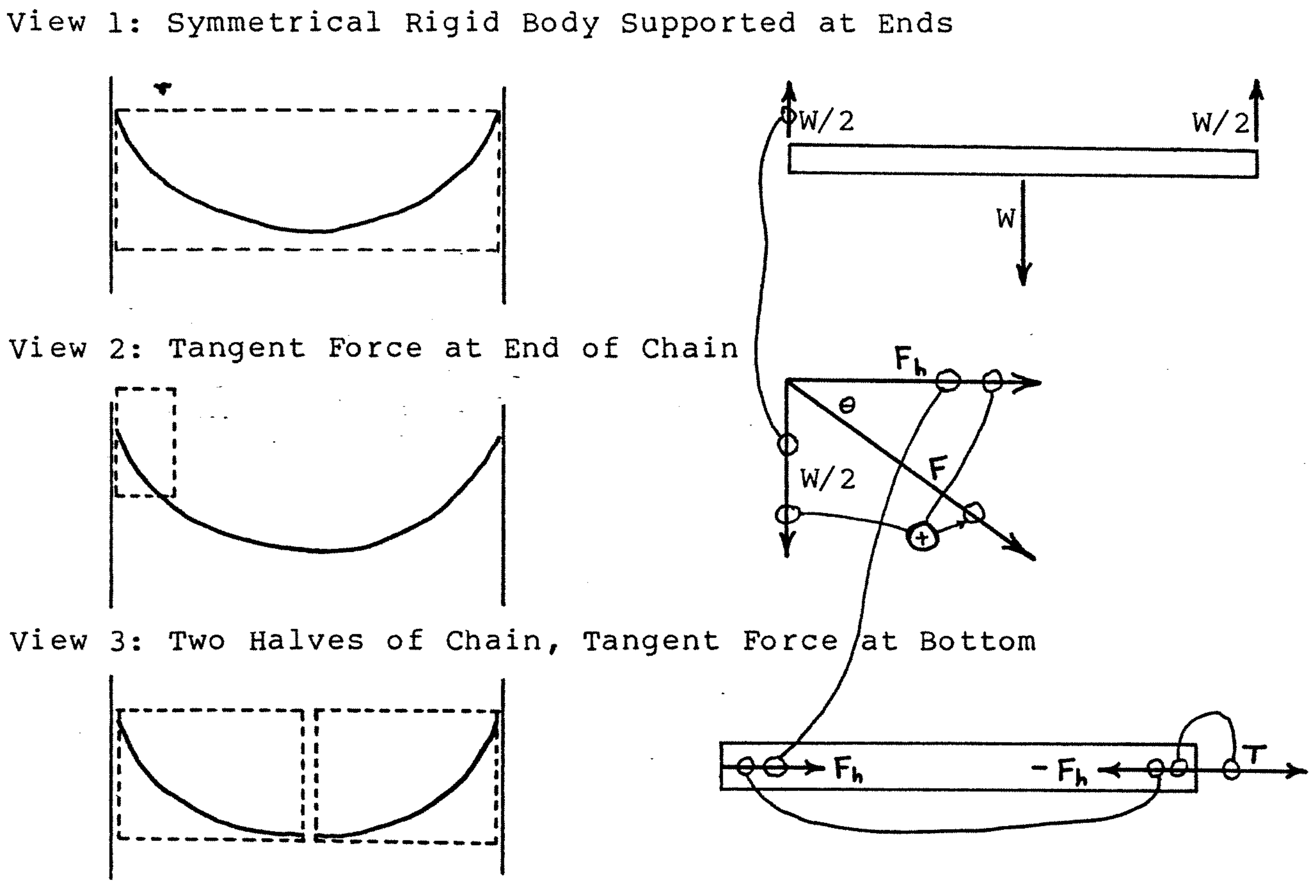

Figure 5: Multiple Views Used in Solving Chain Problem

First, the chain as a whole can be viewed as a rigid body supported at its ends; this is a special case which immediately implies that the vertical force at each support is W/2. Next, we can view one end of the chain; since the chain exerts a force along its tangent and since the vertical component is known, both the total force and the horizontal component can be found. Finally, we view the chain as two rigid bodies attached at the bottom; this is another special case which shows that the tension at the bottom of the chain equals the horizontal component at the top. The problem solver has created the equations needed by taking multiple "gedanken" views of the system and relating these views by identification of their components. The ability to design an analysis method in this manner is the key to truly expert problem solving.

5.Ø CURRENT WORK

We are currently implementing a physics problem solver in accordance with the principles outlined here. We have implemented a representation language (GIRL) and a language to access these representations (GLISP) [15]. The goal of these systems is to allow the specification of a precise internal representation formalism while allowing efficient but informal access to the representations from within LISP programs. For example, one can create a lever object with a statement such as

(A LEVER WITH LENGTH = '(1Ø FT)).

If we wish to view the length of the lever in meters, this can be done with the statement

(VIEW THE LENGTH OF THE LEVER WITH UNIT = 'M).

We hope that this English-like programming will allow us to experiment with different representations without requiring large changes to the programs which access the representations. The analysis models used in solving a problem will be represented explicitly by our system rather than being lost as equations are written.

As we have described it, the process of solving physics problems is more akin to medical diagnosis [16] than it is to theorem proving: we wish to diagnose any of a large number of possible problem syndromes, and prescribe a specific set of analysis methods for each. We plan to use a hierarchical rule-based system (which may also be profitably viewed as a discrimination network or as a generalized parser) to recognize problem types. We believe that if these techniques prove successful in physics, they will be useful in a wide variety of problem domains.

6.0 ACKNOWLEDGMENT

This research was supported by NSF award No. SED-7912803 in the Joint National Institute of Education - National Science Foundation Program of Research on Cognitive Processes and the Structure of Knowledge in Science and Mathematics.

7.0 BIBLIOGRAPHY

1. Halliday, D. and Resnick, R., Physics (Parts 1 and 2), New York: Wiley, 1978.

2. Simon, H. A., The Sciences of the Artificial, M.I.T. Press, 1969.

3. Larkin, J., McDermott, J., Simon, D. P., and Simon, H. A., "Expert and Novice Performance in Solving Physics Problems", Science, vol. 208, no. 4450 (20 June 1980), pp. 1335-1342.

4. Cohen, Harvey A., "The Art of Snaring Dragons", M.I.T. A.I. Lab Memo No. 338, 1975.

5. Sussman, G. J. and Steele, G. L. Jr., "CONSTRAINTS - A Language for Expressing Almost-Hierarchical Descriptions", Artificial Intelligence, vol. 14, no. 1 (August 1980).

6. Bhaskar, R. and Simon, H. A., "Problem Solving in Semantically Rich Domains: An Example from Engineering Thermodynamics", Cognitive Science, vol. 1, no. 2 (April 1977), pp. 193-215.

7. Bundy, A., Byrd, L., Luger, G., Mellish, C. and Palmer, M., "Solving Mechanics Problems Using Meta-Level Inference", Proc. 6th IJCAI, Tokyo, 1979, pp. 1017-1027.

8. van der Merwe, C. W., Schaum's Outline of College Physics, New York: McGraw-Hill, 1961.

9. Leighton, R. and Vogt, R., Exercises in Introductory Physics, to accompany The Feynman Lectures in Physics, Reading, Mass.: Addison-Wesley, 1969.

10. Hofstadter, Douglas R., Goedel, Escher, Bach, New York: Basic Books, 1979.

11. Borning, A., "THINGLAB: A Constraint-Oriented Simulation Laboratory", STAN-CS-746, Computer Science Dept., Stanford University, 1979.

12. Novak, G., "Computer Understanding of Physics Problems Stated in Natural Language", American Journal of Computational Linguistics, Microfiche 53, 1976.

13. Novak, G., "Representations of Knowledge in a Program for Solving Physics Problems", Proc. 5th IJCAI, Cambridge, Mass., 1977, pp. 286-291.

14. Novak, G. and Araya, A., "Research on Expert Problem Solving in Physics", Proc. 1st AAAI Conference, Stanford, CA, Aug. 1980, pp. 178-180.

15. Novak, G., "GIRL and GLISP: An Efficient Representation Language", submitted to 7th IJCAI.

16. Davis, R., Buchanan, B., and Shortliffe, E., "Production Rules as a Representation for a Knowledge-Based Consultation Program", Artificial Intelligence, vol. 8, no. 1 (Feb. 1977), pp. 15-46.

17. Weldon, David, Psychology Dept., Washington University (St. Louis), personal communication.