COGNITIVE SCIENCE 6 37–75 (1982)

Unlearning Aristotelian Physics: A Study of Knowledge-Based Learning*

ANDREA A. DISESSA
Division for Study and Research in Education
Massachusetts Institute of Technology
Cambridge, MA

A study of a group of elementary school students learning to control a computer-implemented Newtonian object reveals a surprisingly uniform and detailed collection of strategies, at the core of which is a robust "Aristotelian" expectation that things should move in the direction they are last pushed. A protocol of an undergraduate dealing with the same situation shows a large overlap with the set of strategies used by the elementary school children and thus a marked lack of influence of classroom physics training on this student's naive physics. The data from these two studies are pooled and elaborated into a "genetic task analysis" of how one might come to understand Newtonian dynamics as a more or less natural evolution from the naive state.

1. INTRODUCTION

For problem solving in domains like physics, no one disputes an important role for prior, domain-specific knowledge. For learning, however, it is easy to overlook the naive knowledge state in favor of a focus on general learning mechanisms, or on the representation and function of knowledge in already competent systems. This paper aims specifically at charting the naive knowledge state and its implications in learning. In particular, we make an analysis of the naive knowledge which students can bring to bear on a standard college curriculum subject (Newton's laws), including some aspects of vector algebra that are necessary to understand the laws.

What we are after is something like a task analysis—what does it take to operationally understand Newton's laws? But given the developmental

*This work supported in part by the National Science Foundation under grant number 77-19083SED, and in part by the Ford Foundation.

*I am grateful to John Seeley Brown and Allan Collins for careful reading of a previous version of this paper which had considerable impact on its final form. I would also like to thank S. Papert, H. Sinclair and J. Richards for insightful comments.

stance expressed above, we intend to base our analysis on empirical data about the naive knowledge state, including some information about paths of local development, and at least a plausible view on how global development could reach "expert" status. Such a *genetic task analysis* will constitute the spine of our results. An important byproduct of our attempt to link naive and expert understanding will be an enriched sense for how intuitive or commonsense knowledge can serve as part of the encoding of "abstract" knowledge, even in experts.

Although the concept of a genetic task analysis is intended to introduce an empirical basis to what is often regarded as an *a priori* activity, there will remain hypothetical elements, perhaps necessarily so. In the first instance, in order to make the project of finding a genetic task analysis experimentally tractable at all, one must circumscribe the pool of naive knowledge, and the routes of access to it which will be considered. We will focus on the knowledge that students spontaneously apply to understand and control a particular computer-implemented environment. The environment involves a simulated object called a *dynaturtle* whose motion is directed by student commands in accordance with Newton's laws.

A second hypothetical element inherent in making a genetic task analysis stems from the attempt to establish a path of conceptual development, or at least a class of paths. The present study taps into only a local development (one hour to two weeks), and hence, we will need to extrapolate in order to imagine, even roughly, the complete naive/expert transition. More profoundly, one can never rule out the possibility of radically different routes of development based perhaps on entirely different pools of naive knowledge. Neither of these elements of hypothesis can be simply addressed, and we will need to return to them, in various guises, throughout the course of the paper.

The organization of the paper is as follows: Section 2 introduces the computer environment used for the study and gives a naturalistic account of the behaviors of a set of elementary school students encountering it. The results indicate a surprising structure of discrete and definite "theories" rather than a continuous acquisition of skill in controlling the dynaturtle. At the core of this naive physics is a robust Aristotelian expectation that objects should go in the direction they are pushed. The richness and robustness of this set of theories we take as important *post hoc* justification for the particular choice of ways of tapping naive knowledge relevant to understanding physics.

Section 3 abstracts these results into a "learning paths chart," the aim of which is to catch the essential features of development seen in the students. Section 4 extrapolates the empirical data into a more complete genetic task analysis. Section 5 turns to the question of what such analysis, based on very young students, can have to say about learning physics in the university. We give a synopsis of a protocol of a freshman college student exposed to the dynaturtle environment after a year of high school physics, and nearly a term of college physics. Her behaviors can be matched in large degrees to elements of the learning paths chart. This snapshot of conceptual development at a much later stage provides important confirmation of the importance of considering the naive knowledge tapped by dynaturtle in the elementary school students. It also indicates stages of development beyond those accessible to the younger students, and suggests a functional role in expert understanding for the refinement of the naive physics which students developed to cope with the dynaturtle.

2. NAIVE PHYSICS IN GRADE SCHOOL

The empirical basis for this section is a set of naturalistic observations made of eight sixth grade students (two classes of four) in a computer laboratory located in a school in suburban Boston (Papert et al., 1979). The students had been chosen to represent the entire range of academic abilities, and they had previously had eight weeks of roughly four hours per week experience using the computer language Logo in a classroom environment which encouraged each child to select his own projects and activities. During the final two weeks (total 10-week Logo exposure), the dynaturtle was introduced as another domain of activity for the students. Most of the students chose to devote a substantial proportion of their remaining time to dynaturtle. (Two students were deliberately encouraged to stay with previous work and did not figure in the study.) Observations were made by an in-class observer and the class's regular teacher.

2.1 Description of Dynaturtle

A dynaturtle, like its ancestor the Logo geometry turtle, is a graphics entity which can be moved around on a CRT with commands typed at a keyboard. The geometry turtle formed the basis for most of the work done previously with Logo by the children in the study. Both turtles respond to commands, RIGHT or LEFT by instantly turning in place. A number following the turn command, called its input, tells how many degrees to turn. Translational motion for the geometry turtle is caused by the command FORWARD. FORWARD 100 causes the turtle to move 100 steps in the direction it is facing. In contrast, a dynaturtle never changes position instantly, but can acquire a velocity with a KICK command which gives it an impulse in the direction the dynaturtle is currently facing. To effect real time control, one normally directs a dynaturtle with single keystroke commands, R, L, and K which stand for RIGHT 30, LEFT 30 and KICK 30.

The dynaturtle represents Newton's laws in the following ways: (1) Newton's first law is that an object must remain at rest or travel at a constant speed in a straight line when no force is acting on it. This is precisely the behavior of the dynaturtle when no kicks are applied. (2) Newton's second law, symbolically $F = ma$, specifies how velocity changes when a force is applied: the vector change in velocity is proportional to the force and inversely proportional to the mass of the object. A dynaturtle kick is a discrete (impulse) version of a force and specifies the change of velocity according to the discrete version of $F = ma$ which is depicted in Figure 1. (The direction of the kick vector is the heading of the dynaturtle, and the magnitude is proportional to the input to KICK, or the number of kicks in real time mode.) The figure indicates the central role that vector addition plays in understanding Newton's laws. Since only one object is visible, the effect of different mass on the result of a kick is not functionally modeled. Neither is Newton's third law (action and reaction) represented, which prescribes the effect of a kick on the agent (kicker).

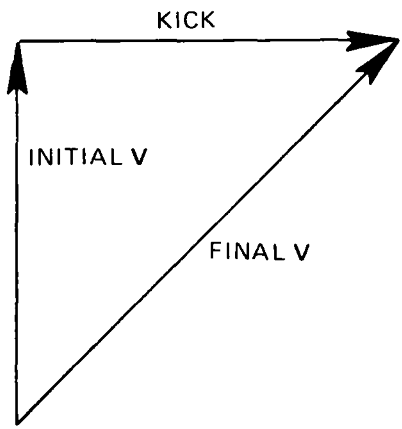

Figure 1. Discrete version of F=ma.

Two model games were provided for the students. The relevant one here was called Target. Its goal was simply to direct a dynaturtle with K's, R's, and L's to hit a target, but to do so with a minimum speed at impact.[1] A qualitative scoring, e.g., "too fast," together with impact speed was printed out when the target was reached. The initial configuration had the dynaturtle at rest aimed directly up the screen and the target, as indicated in Figure 2, positioned at bearing 45° from the dynaturtle. A single K command would cause dynaturtle to travel the distance between initial position and target in about 15 sec. The introduction to dynaturtle given to students was a brief description of commands together with an illustration, applying a few "kicks" to a tennis ball on a table using a small wooden mallet.

[1]Even with fixed magnitude kicks and turns, it is possible to hit the target with an arbitrarily small speed.

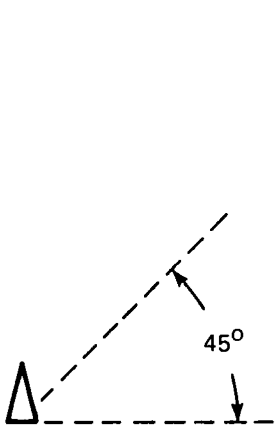

Figure 2. Initial configuration of Target game.

2.2 Overview of Results

In view of the striking differences in abilities and style which the students exhibited in their other work, we were greatly surprised to see how uniform their responses were to the dynaturtle. Students seemed to have definite non-Newtonian expectations which were contradicted by the behavior of the dynaturtle. In fact, one might characterize early stages of students' work as the confrontation of an essentially Aristotelian theory of physics with a Newtonian reality. For our purposes, we use the term "Aristotelian physics," to mean that objects simply move in the direction you push them. The conflict is whether force correlates with changes in velocity (Newton) or with changes in position (Aristotle).[2]

The germ of the conflict resides in a simple situation which all the students encountered and all regarded as problematic. Suppose a dynaturtle is moving upward, and one wants it to move to the right (Figure 3a). The Aristotelian strategy is simply to aim to the right, then kick in that direction. The expectation is as shown in Figure 3b. "Kick to the right means move to

[2]In using the term Aristotelian physics, we mean generally to impute a definite but non-Newtonian stance to our subjects. More specifically, Aristotle's theory of "violent" (forced) motions is very close to the expectations exhibited by our subjects, specifically with respect to the lack of concern for the effect of previous motion in predicting the results of a force. Aristotle's image for force of this kind was "carrying" which left no room for motion without, or independent of a force. He thus had to invent an *ad hoc* mechanism to explain the effect of momentum in keeping an object moving, and never attempted any principle of combining that effect with another force. For different views on what "Aristotelian physics" should mean with respect to naive understandings of $F = ma$, see Cohen (1974) and Shanon (1976); diSessa (1978) gives an early version of the view expressed here.

the right." In contrast, the Newtonian dynaturtle moving upward has momentum in the upward direction which is not affected by the sideways kick, and thus it takes a "compromise" path away from the kick as shown in Figure 3c. The vector addition relevant to this is that shown in Figure 1. All of the students spontaneously generated the sideways kick as a means of making a right turn, and expressed surprise and consternation at the result. Complaints that the machine was not working correctly at this point were commonplace and vociferously made. The robustness of the students' theory is attested to by the fact that, though many of the students had made significant progress in the two weeks of exposure, none proved to completely shed the Aristotelian disposition. Equally intriguing and more important from the standpoint of making a genetic task analysis is that despite an incorrect "theory," students proved capable of developing alternate strategies for dealing with the corner situation, strategies based on other ideas more compatible with Newtonian dynamics. In fact, several achieved practical mastery of the dynaturtle in most circumstances and proceeded to use it in projects of their own.

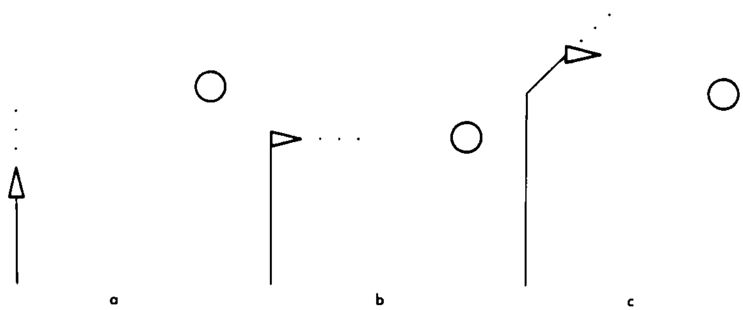

Figure 3. (a) Moving upward; (b) expectation after sideways kick; (c) actual result.

The following gives a more detailed account of the activities of two of the students who were observed at greatest length.

2.3 Detail

Jack's early work with the dynaturtle was typical of the group in the sequence of strategies, successes and failures. His initial plan with the target game was a simple *Aim and Shoot*[3], the almost universal starting place of all subjects:

[3]Aim and Shoot is one of the principal strategies involved in playing with the geometry turtle, and it may therefore be a straightforward transfer to find it immediately and clearly implanted in this slightly different environment. On the other hand, informal observation of turtle-naive adults has also shown this to be a near universal starting point. No turtle-naive children have been observed.

1. Turn the turtle with R's and L's until it is facing the target.
2. Shoot using K's.

This plan cannot succeed as it stands: the target is at a 45° bearing, and R and L make 30° increments, thus Aim and Shoot necessarily carries the dynaturtle off to the left or right of the target. Once Jack saw the failure, he summarily dropped the strategy. Other students modified the plan to accept 30° or 60° as close enough for a start. There is every reason to believe, especially considering his experience with the geometry turtle, that Jack understood the problem and simply looked for alternatives. This is in sharp contrast to what happened with his next strategy.

The alternative plan Jack (and most others) adopted was to move straight up the screen, then, when the dynaturtle was at the same height as the target, make a right hand 90° turn and run into the target by kicking toward it.[4] Of course, this *Aristotelian Corner* strategy brought Jack quickly to the heart of the Aristotelian-Newtonian controversy. The turtle skipped diagonally away from the target rather than toward it. His first instinct on failing was to try it again and again. Then, he applied more kicks at the right angle turning point. Aim and Shoot had failed for an understandable reason —he did not complain when it did. But, his attempted Corner had no good reason for failure in his eyes, and he complained and appeared frustrated.

At this point, an intervention was made, discussing with him the essential difference between turtle and dynaturtle. Out of the discussion arose a new strategy which was neither explicitly proposed to him, nor entirely spontaneous on his part: at the corner, stop the turtle with kicks in the opposite direction of its motion, then Aim and Shoot directly into the target. (The stopping kicks which canceled initial kicks were named *antikicks* by another student who independently and spontaneously proposed the idea. We appropriate the name.) Jack understood and quickly adopted this strategy, which we will call a *Newtonian Corner* (Figure 4).

There are two significant points to make at this juncture. Jack never did exhibit any confusion between turning and kicking; that they are independent actions (perhaps modeled on the independence of move and turn commands for the geometry turtle) was taken for granted. For example, in trying the Aristotelian Corner, he turned the 90° immediately after kicking to start the turtle, then waited patiently to give the second kick when the tur-

[4]It is possible that the universality of this step with the children is due to the fact that they have already had significant experience with the geometry turtle. This corner movement is a very frequently observed strategy in that domain. It is used to achieve accurate positioning (as in positioning parts of a picture). On the other hand, the strategy had been neither taught nor even named or remarked upon. Naive adults have been nearly as uniform as the children in applying this strategy to Target.

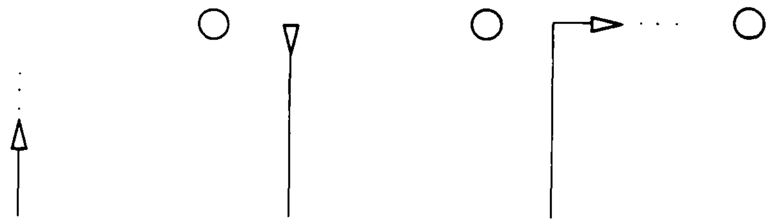

Figure 4. (a) Kick to start; (b) turn and kick to stop; (c) turn and kick to finish.

tle reached the corner. This is important, as it shows that he did *not* have trouble disassociating aiming from moving in his switch from geometry turtle to dynaturtle. Without this fact, one would be tempted to attribute Aristotelian expectations to a simple carry-over of the fusion of direction-pointed and direction-moved which characterizes the geometry turtle. The lack of difficulty in differentiating the direction of motion and direction of pointing was true of the other students as well.

Secondly, Jack knew without being told and before experimenting that the number of kicks he needed to give to stop the turtle was the same as the number he gave to start it. He did not worry about timing but only about number.

Having developed a foolproof strategy which he understood, Jack concentrated for an extended time, practicing and elaborating it.

Donna started out with the same Aim and Shoot as Jack but was more patient in trying to debug it. She accepted 30° bearing as close enough to start and, because of her care not to give too many kicks, she in fact succeeded in hitting the target, but not reliably. Trying to follow Jack's corner path, (she could see his screen and what he had done, but apparently did not see how) she fell into the same, Aristotelian trap. Again at this point, an intervention was made to assure her that despite her complaints to the contrary, the computer was working properly. Side kicks and the resultant diagonal trajectory were illustrated with the tennis ball and mallet. This appeared to engender a state of disequilibrium. She made it clear with facial expression that she was quite dubious about this "experiment," grabbing the mallet and trying it herself several times. "There must be a way," she said, continuing to try versions of the experiment, for example, twisting the mallet as she hit the ball. (We are not sure what she intended by her twisting hits, except to try to make the sideways kick work as she expected.) She was shown Jack's Newtonian Corner strategy of a hit to stop, reaim and new hit, but she still indicated a wish to see the corner accomplished with one kick. A diagonal backward kick (at bearing 135° to the velocity) was suggested and demonstrated, but she refused even to consider that. "I like

Jack's strategy," whereupon she returned to try it out on the computer. Such a diagonal backward kick is the equivalent to centrifugal force in the discrete dynaturtle world, and is an important but counter-intuitive Newtonian strategy.

Donna was not content as Jack was to stay with one method. Over the next few days, she tried many others. In particular, she tried the Aristotelian Corner strategy starting out *horizontally* rather than vertically. Of course, it failed. She tried to correct a 30° attempt at Aim and Shoot by kicking perpendicular to the established trajectory when that aim ran directly through the center of the target. It failed as well. She tried and failed at correcting the original Aristotelian Corner's defects with yet another kick, but using the same "kick toward the target" strategy. It is not clear how to interpret this insistent repetition. On the one hand, one might say it simply reflects the strength of her Aristotelian convictions. However, the analysis developed below suggests a potentially crucial developmental function.

The crucial point is that understanding these attempts as repetitions requires a way of thinking about them as being the same. Though from an adult perspective, the fact that they are the same might appear evident, that is not necessarily so for school children. Some Logo students of this age will not recognize a "diamond" with four equal and mutually perpendicular sides as a square (produceable with a "square" turtle program), presumably because orientation is relevant to their encoding of shape. Concerning dynamics, the situation is further complicated by the fact that in the real world, gravity breaks the underlying rotational symmetry. For all Donna knew from her initial experiment, moving up and kicking horizontally might be a singular case.

If we pretend, for the sake of discussion, that Donna's internal representation is verbal, her initial description of the Aristotelian Corner might have started out something like, "It doesn't go straight when you kick it sideways," or even, "Sometimes it doesn't go straight." Such descriptions offer only a vague hint of what is actually involved.

Seen in this way, Donna's experiments make utterly clear scientific sense as a way of developing and refining a mechanism for understanding the identity of those separate events. One possible mechanism, again assuming verbal representation, would be to use invariant language, that is to say, language in which the description is unchanged if a different point of view on the experiment is taken. Such language assumes that the frame of reference (up, down etc.) has no functional significance. "Kicking perpendicular to an established trajectory" is invariant in this sense. Of course, there are other mechanisms, and it is not likely that any such development would be purely verbal. But, despite the lack of direct evidence or detail, the main point remains: It is almost certain that many students of this age need to develop an invariant recognition capability of some sort to see the Aristotelian Corner and other dynamical phenomena as "things." This is not trivial learning about physics, a field in which symmetry is profoundly fundamental; and, it also seems clear that repetition is a requisite part of any untutored learning of it.

Donna was presented another workable refinement of the Aristotelian Corner strategy, other than Jack's start and stop Newtonian Corner: After starting upward, "cut the corner," turn, and kick sideways early. This particular idea, the *Early Strategy* (Figure 5), was less common than antikick as a spontaneous notion. However, it makes good intuitive sense; thinking of dynaturtle as being slow to respond is a good heuristic, and when proposed, it was readily accepted.

Donna instantly adopted the Early strategy, and spontaneously added the corrective feedback loop, "If you miss by getting to the target too late, (meaning x-coordinate position "gets to the target" after y-coordinate) then kick earlier," and vice versa. Kicking "too late" or "earlier" in fact are expressions universally used by our subjects to describe the phenomena, and their own intent. It doesn't seem difficult to understand the naturalness of such a strategy. For example, in getting to school at a certain time, one may employ the early/late conceptualization and feedback loop with respect to the question of when to leave home. On the other hand, the terms in which we are forced to phrase our description of this idea as applied to this situation involve fictitious events, crossing one or another coordinate, and are surprisingly abstract.

After quite a bit of play (much of which has been described above), another attempt was made to bring Donna to understand the diagonal backward kick method for making a Corner. She was asked to think of using the Newtonian Corner method (two perpendicular kicks: one to stop, and another kick toward target) but with the kicks coming very close together in

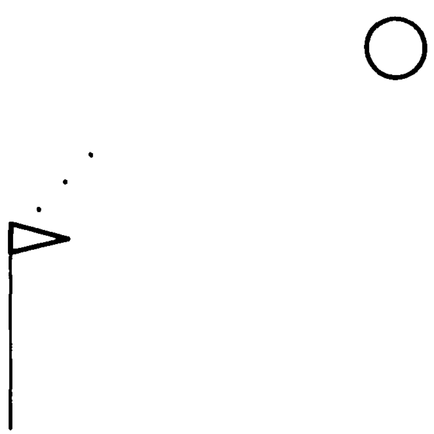

Figure 5. Early strategy.

time. Now think of a single kick having the same effect as the two. "I know" she said. "You want to kick at 45°!" She meant a 135°, diagonal backward kick, as became evident at the computer. She was in fact anxious at this stage to try out the method. (Though 135° is an unattainable turn, two kicks at 120° does turn the corner.) This is a solid leap. She is making explicit a qualitative version of vector addition in proposing for herself the diagonal backward Corner strategy which, before experimentation, she had emphatically rejected as impossible or incomprehensible.

3. A LEARNING PATHS CHART FOR TARGET

We would like to abstract and condense the data from observations like the preceding. To first approximation, we want a list of student theories and hypotheses about the dynaturtle. We shall state these as strategies or "issues of concern" which the students evidenced, without looking for a deeper structuring. Generically, elements of such a list will be called *topics* metaphorically to imply that we can read a student's protocol as a series of topics which expresses aspects of the student's growing understanding of the Newtonian dynaturtle. It will be natural to include comments on topics such as success or failure (of strategies) and students' reactions, such as surprise or matter-of-fact attitudes.

More importantly, one would like to include information on possible development. This cannot take the form of a simple sequencing since many variations in the order of topics were observed. (As is common sense about human nature, there is a profound instability in people's approaches to complex domains; exactly when a particular idea will occur is difficult to say, even if it is certain to occur.) Instead, we will draw directed links between topics which indicate that such a topic-to-topic transition was observed. Most of the links can be understood as refinement or debugging of already known strategies. Some indicate tighter relations (two-way arrows); for example, though antikick appears to be a prerequisite substrategy of the Newtonian Corner, it seemed almost always to emerge out of the need or wish for a stopped state in order to accomplish the corner movement. Some judgement was involved in excluding links which were considered accidental. We do not attempt to understand or classify transitions here, though this would seem a profitable line of inquiry as a way of getting at, for example, prerequisite topics and other necessities of ordering.

Chart 1 is the topic list with developmental links, which we call a *learning paths chart* for the Target game. We would like to think that the behavior of any subject could be traced as a "wandering around" from topic to topic via the noted links on some future version of such a chart. Generally speaking, down the page indicates later in development. Section 3.2 below contains detailed descriptions of all the topics.

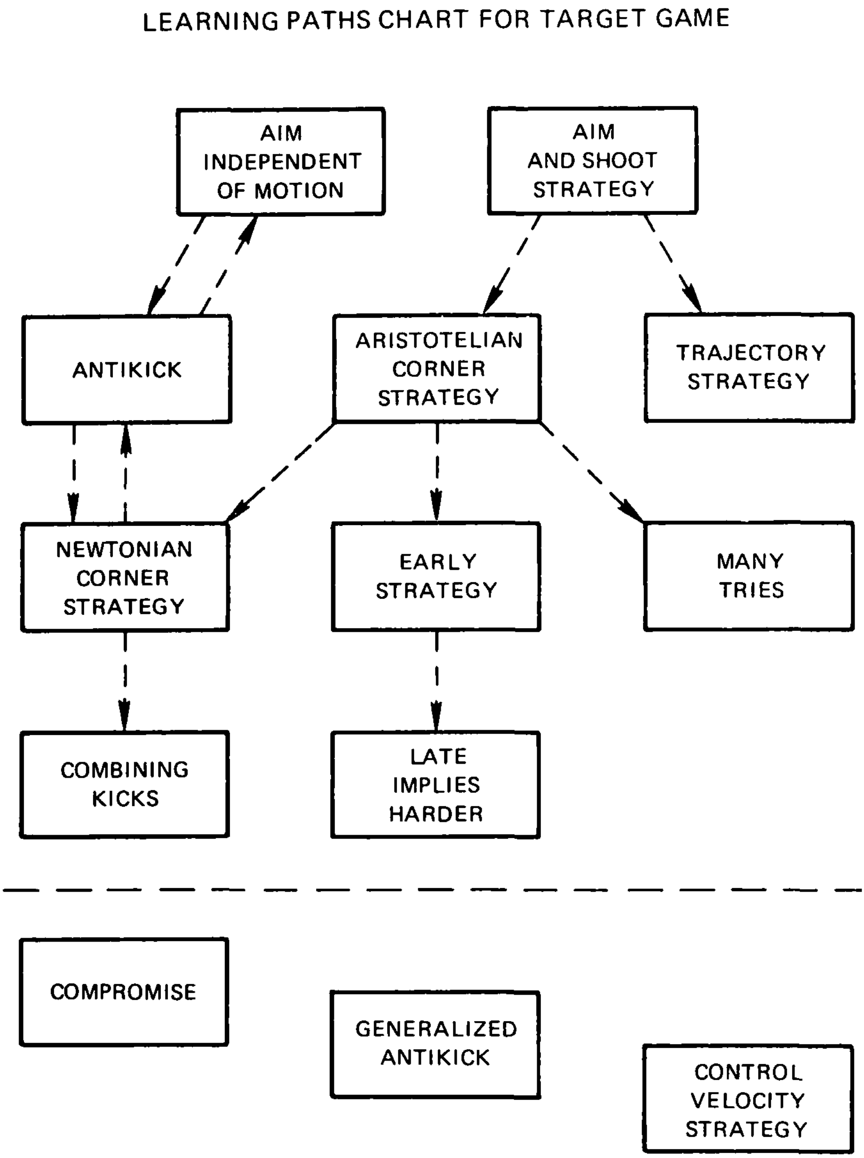
LEARNING PATHS CHART FOR TARGET GAME
AIM INDEPENDENT OF MOTION
AIM AND SHOOT STRATEGY
ANTIKICK
ARISTOTELIAN CORNER STRATEGY
TRAJECTORY STRATEGY
NEWTONIAN CORNER STRATEGY
EARLY STRATEGY
MANY TRIES
COMBINING KICKS
LATE IMPLIES HARDER
COMPROMISE
GENERALIZED ANTIKICK
CONTROL VELOCITY STRATEGY

CHART 1

3.1 When Almost Everybody Does Essentially the Same Thing

The ten topics above the dotted line in Chart 1 were abstracted purely from observation of the school children. We have already remarked that there was a surprising degree of commonality among the students. Subsequent to that study, we have informally observed a collection of physics-naive adults (12 in number) at the same task. Strikingly, the older subjects showed considerable overlap at early stages with children's topics, but seemed to have the capacity to develop faster/farther: to wit, time compression of two weeks to roughly one hour, and three additional topics. These findings are especially provocative considering Piaget's interpretation of similar transitions of strategies as tied to global developmental shifts in the pattern of reasoning. In view of the informality of the data, however, we refrain from speculation aside from remarking that such data increases our confidence that the learning paths chart of Target has tapped a fundamental epistemological issue. We have not attempted to expand developmental links from the children's topics into the adults'.

3.2 Annotation of the Chart

Below are descriptions of the topics appearing in Chart 1. Readers not interested in details may safely skip them.

TOPICS FOR DYNATURTLE TARGET GAME

Aim and Shoot Strategy—"Turn the turtle toward the target and kick" is the most universal starting point. Its failure, due to size of turns incommensurate with 45° bearing of the target, was accepted and never problematic.

Aristotelian Corner Strategy—This is the classic non-Newtonian strategy described earlier, involving the assumption that the turtle travels in the direction of kick. Its failure always resulted in great surprise, repeat tries, and sometimes extended consternation. Note that while one might interpret Aim and Shoot and the Aristotelian Corner as both expressions of the same theory, they are included as separate topics in the chart. This is done because their different contexts provide strongly different strategic outcomes. One satisfies expectations; one does not. We shall have much more to say about context in later sections.

Trajectory Strategy—This strange strategy did not occur frequently, but often enough to warrant including. To debug Aim and Shoot, some seemed to posit a curved path approach; and, as a mechanism for obtaining such, a systematic and repeated pattern of K's and R's were used, e.g. K R K R K R. "Curving" and "starting to curve" were frequent verbal accompanying descriptions. There seemed to be the assumption of a simple relation between the turn-kick combinations and the amount of curving. In any case, the rate and pattern of keystrokes were the parameters varied in an attempt to refine this strategy into a working one.

Aiming Independent of Motion—The default assumption made naturally by everyone was that turning would not affect the direction of motion until a kick was given. Despite this, various circumstances called this hypothesis explicitly into question. The most important of these circumstances was the context of trying out the antikick idea (perhaps worrying about a presumed "minor" effect interfering with the exact cancellation of kicks). Despite Aiming Independent of Motion, students frequently associated reaiming with the subsequent kick, and did reaiming just before the kick, even if there was much empty wait time preceding the reaim-kick combination.

Antikick—The "kicking the opposite way to cancel a kick" phenomenon was a spontaneous idea in most cases, and an immediately accepted suggestion in the rest. Its importance probably lies in the function of achieving the stopped state, and in its intuitive roots in the powerful idea of canceling. Three subspecifications are necessary: (1) Kick-antikick starts and ends at rest; superposition on an initial velocity is not conceptually possible at this stage. (2) It is assumed that any number of kicks will be exactly canceled by the same number of Antikicks, the timing of the kicks not being an issue. (3) Interspersing kicks in other directions between kicks and antikicks disqualified the antikick strategy until a much later stage (see Generalized Antikick).

Newtonian Corner Strategy—The canonical debug of the Aristotelian Corner, involving antikick to stop, reaim and kick toward the target, was almost always associated with right angles (as the Aristotelian Corner was).

The Early Strategy—(Figure 5) This appears to be a more sophisticated strategy than the above. It occurred later in adult protocols and not at all in many of the children's protocols. This strategy involves qualitative reasoning about control of the relative timing of events (e.g. crossing the vertical line containing the target) with respect to another event (reaching the height of the target). It has a distinctly different perspective than the static geometry of the Aristotelian "aim toward the target." Emphasizing this is the fact that no subject engaged in Early considered to think about or observe geometric features, such as the orientation of the path produced by the pair of orthogonal kicks which is, of course, 45°. Had that observation been made, one would conclude that the sideways kick should be given immediately after the start kick in order to follow the 45° bearing to the target. This phenomenon, attending only to features directly relevant to current strategy, occurred in several other instances. See "game 3" described in the appendix for another example.

Late Implies Harder—This refinement of the Early Strategy, "if you are late, you should kick harder (more)," notably did not occur simultaneously with the principal Early Strategy.

Many Tries (Developing Invariant Recognition Capability)—The overt trying over and over of the Aristotelian Corner Strategy in many orientations, as carried out by Donna, was not as frequent as it was striking when it did occur. As if to suggest they had more fluid invariant recognition mechanisms, adults were much less prone to such experimentation. Though other strategies such as Trajectory and Newtonian Corner were also repeated, we have selected out the repetition of Aristotle as a topic because of its hypothetical function in developing invariant descriptive techniques. Other repetitions had more obvious and less interesting reasons; for example, the strategy repeated was the only working strategy the student had, or at least his current best hope for a strategy refinable into a working one.

Combining Kicks Thought Experiment—(See description at the end of Section 2). This is included not because it was universal, in fact it only occurred once with the children, but because it marks a striking advance over initial inclinations. Since, for the most part, we retained a non-interventionist stance, we did not get good data on when and how reliably receptive students are to this idea.

The following three "adult" topics were never observed in the children's work, but have been seen in the work with older subjects. We have good reason to believe that these three topics are not simply more of the same, but instead might reflect an essential change of thinking in coming to appreciate, in a structured way, the vector notion of velocity and its relation to kicks. Previous topics focused on geometric configuration in space (Aristotle, etc.) or time (Early), or on intervention history (kick-antikick) in order to decide what to do. In contrast, Newtonian physics focuses on velocity as the determinant factor of future trajectory. In particular, other than the position, which is seen very early on as a factor, the only thing needed to summarize all previous interventions is velocity. Thus, coming to treat velocity as a central entity relevant to choosing intervention appears a major transition from naive physics toward Newtonian physics. In the following descriptions of strategies, we will comment on their relations to coming to understand the nature of vector velocity as a summary of previous interventions and predictor of the effect of future ones. To keep these two classes of topics separate, we will call the earlier (elementary school) set—*intervention history* topics—and the later ones—*velocity-centered* topics. The names serve to remind us of the possibility of a genuine stage transition between them. In any case, the remote appearance of velocity-centered strategies, especially compared to antikick, would seem to indicate the difficulty of understanding dynaturtle on the basis of its *state* summarized in its velocity, independent of intervention history. This contrasts strongly with the very early and often covert assumption in physics texts that position and velocity are sufficient initial conditions, i.e. give a sufficient summary of the entire history of the particle to predict all future motion.

VELOCITY-CENTERED TOPICS

Generalized Antikick—This signifies the realization that any sequence of kicks in any set of directions can be selectively and exactly undone with antikicks on a one-for-one basis, and in fact, in any order. Very likely, this topic will be refined by further study into stages involving, for example, noninterference of perpendicular kicks as a first noticed special case. The Generalized Antikick phenomenon was employed usually to obtain zero velocity or to explain an accidental zero velocity state. It is an advance over kick-antikick in that it involves the implicit recognition that the effect of a kick remains after, and in fact is unaffected by subsequent kicks. Thus, it can be "undone" at any time with an appropriate antikick. From a Newtonian perspective, of course, velocity is the repository of previous kicks and by virtue of being a vector sum of kicks allows selective canceling with antikicks.

Compromise—Combining Kicks involved a qualitative composition of a pair of kicks to form an equivalent single kick "between the two in direction." A formally similar operation combines an established velocity with a kick to produce a new velocity in a direction between the two. This we call Compromise. We saw little indication of Compromise in school children. (Though, we did not try to elicit it, either, as we did Combining Kicks.) Again, the reason Compromise appears to be a more advanced notion than Combining Kicks is that it involves the recognition of the existence of velocity (momentum might be more appropriate —the tendency to keep going as is), and also because the operation combines two different kinds of things (velocity and kick).

Control Velocity Strategy—This topic finally represents the use of velocity as an intermediate construct summing up the effect of all previous kicks. It embraces two sub-strategies: (1) Kicking opposite velocity reduces speed. (Kicking to increase velocity seems to have been used much earlier.) (2) Kicking perpendicular to velocity changes direction. These two can be used in combination to achieve both effects simultaneously with kicks at various angles, with respect to current velocity.[5]

[5]It might be useful to mention two relatively simple strategies which, somewhat surprisingly, were never observed with adults or children. Despite the frequent use of the Newtonian Corner (Jack, for example, did very little else), the following simple modification never occurred. One can save some time with no penalty by "cutting the corner;" simply apply the sideways kick to the dynaturtle before giving the antikick that stops the forward component of motion. The second unobserved strategy involved the frequent occurrence that, after a large number of kicks, subjects accidentally wound up with a small velocity, but not directed at the target. Unfortunately, no one could take advantage of that through the realization that that small velocity could be recovered at any future time by antikicking any subsequent kicks. One need only use those subsequent kicks to "reposition" the small velocity, so that, after it is recovered, it intercepts the target.

4. A GENETIC TASK ANALYSIS

In principle, we should have no *a priori* reason to suspect that the naive strategies for dealing with a situation, such as the target game, could do more than suggest why students have difficulty with Newton's laws. What seems special in this case is the remarkable richness and potency of those strategies. In fact, this section's aim is to elaborate a genetic task analysis of the Newtonian mechanics embodied in dynaturtle based only on the topics appearing in the learning paths chart. We will aim at producing a curriculum-like sequence of topics, essentially just a sequence of topics from the learning paths chart, which together cover the subject. Though we will not be precise about what "covering the subject" means, we believe what follows can easily be refined in that way. The analysis will be in two stages: what does one learn from dynaturtle; and how might that evolve toward expert physicist understanding?

4.1 Controlling Dynaturtle viewed as Physics

Claiming that learning to control a dynaturtle is in itself learning some important physics may be controversial; it violates some accepted conventions about what physics is. Though such a claim will be bolstered by the analysis in later sections, we take a more modest view here, only to propose learning to control dynaturtle as an analogy to learning physics which can tell us something about the "real thing." Bear in mind, however, that the elementary school students are learning to deal with (simulated) physical phenomena in a way in which they could not initially, even though that contact is not mediated by symbols and conventional formalisms. Recall also the suggestion that students might be acquiring the ability to focus on features relevant to invariant description, e.g. the relative bearing of kick to initial velocity. Finally, recall that dynaturtle was designed to mirror at least 1½ of Newton's 3 laws.

What kind of perspective does the learning paths chart put on learning to control a dynaturtle? For contrast, let's initially consider a more conventional task analysis by first looking at dynaturtle from the viewpoint of an expert physicist. Try to put aside one of the important experimental results implicit in what has been said so far, that average sixth grade students *can* learn to drive dynaturtles.

A parsimonious description of dynaturtle can involve a vector component of state (velocity) and a state changer (kick) which increments velocity by vector addition. The task analysis might quite reasonably begin with the notion of instantaneous velocity and vectors, including vector addition and component decomposition. Thus, we see appearing a familiar list of prerequisite studies (for college students!) including perhaps analytic geometry, trigonometry, and some fundamentals of calculus. Can we seriously think a fifth or sixth grade student can understand dynaturtles and manipulate them without months or years of study, or "at best" by learning by rote incomprehensible algorithms?

Now contrast the following task analysis based on the Target learning paths chart which we summarize in the form of a sequence of natural (and in some cases trivial) abstractions of experience with dynaturtle.

0. Establishing an invariant recognition capacity of dynamical phenomena, as Donna seemed to be doing in rehearsing the Aristotelian Corner (Many Tries).
1. The remark that aiming does not affect motion (Aim Independent of Motion).
2. The warning that Aim and Shoot fails when the dynaturtle is in motion (Aristotle).
3. The phenomenon of Antikick and its powerful use in producing a true Newtonian Corner strategy.
4. The Early strategy and its refinement, Late Implies Harder.
5. The thought experiment of Combining Kicks, as at the Newtonian Corner (or the reverse, thinking of a diagonal kick as a backward kick to cancel present motion plus a sideways kick to establish a new direction).
6. Compromise, qualitatively combining momentum and kick.
7. Generalized Antikick, that is, ignoring the potential interference of intermediate kicks or an initial velocity in the antikick canceling mechanism.
8. The strategy of Controlling Velocity.

This list is essentially a path through the learning paths chart which happens to be both a sort of "average" observed path and a seemingly natural pedagogical path. For reference, we will call this particular path the *modal path.*[6]

We must add several caveats to our description of the modal path. Topic 0 is a somewhat different order of learning, and does not have as natural a "place" in the sequence as the other topics. Recall that 1 through 5, intervention history topics, were empirically seen with elementary students. But the status of 6 through 8, velocity-centered topics, is less certain,

[6]More recent work (McDermott *et al.*, 1980) gives greater statistical weight to the claim that the modal path, particularly the sequence Aristotelian Corner → Newtonian Corner → Early → Late Implies Harder, is a reliable expectation for spontaneous behavior at the college level.

as regards physics-naive persons. We are not certain that elementary school students could understand these elements, nor even that they represent possible results of unschooled learning.

4.2 Physics Beyond Controlling Dynaturtle

Now we extrapolate toward expert understanding. A small, further abstraction brings the modal path much closer to recognizable, textbook physics. 0 has already been abstracted to this level in our previous description of it. Numbers 1 and 2 combined are a strong affirmation that force and direction of motion are uncoupled. Number 3 proposes a very special case of vector addition, $v - v = 0$ or $v + v + \ldots + v - v - v - \ldots - v = 0$, which in the context of action and "undoing" counteraction seems very intuitive. The idea is restricted in that the vectors added are semantically kick actions and cannot represent velocity. It is also important to note that this intuition is lost, if confounded by intervening kicks, or even if it is superimposed on an established velocity. In number 4, the Early strategy represents an important step forward in that it accepts, and even uses, to good advantage, the initially counterintuitive diagonal motion resulting from the Aristotelian Corner. It and especially its refinement, Late Implies Harder, involve qualitative versions of vector addition, in this case vector addition of an established velocity with an impulse. It is especially nice that students naturally did this with right angle kicks so that the pedagogically important special case, orthogonal composition, is exercised. In number 5, combining kicks at the Newtonian Corner is another step toward understanding the full implications of vector addition. In this case, two kicks are added to each other to produce a theoretically equivalent kick. Again, the natural right angle context is pedagogically advantageous. In number 6, compromise extracts the qualitative vector addition of Early and its refinement from the irrelevant (from the point of view of expert physics) dynamical feedback loop of conceptualizations. It represents another step along the road to recognizing velocity as a vector quantity to which another semantic category, kick, can meaningfully be added. In number 7, Generalized Antikicks is an important special case of commutativity and associativity of vector addition, $(v_1 + v_2 + \ldots + v + \ldots + v_n) + (-v) = (v_1 + v_2 + \ldots + v_n)$. In 8, Control Velocity represents a most advanced stage of understanding the general control of a Newtonian object via pushing around, with impulses, the intermediary state abstraction—velocity.

This analysis has proposed a detailed refinement of what, in a standard curriculum, might go into a monolithic chunk entitled vector addition (of velocities, forces, and impulses). It is refined in at least two important respects. First, it singles out particularly easy and particularly difficult special cases, both of pedagogical interest. It specifically points out a pre-eminent bug in naive physics, the Aristotelian expectation. But more than that, it points out particular situations where students spontaneously do the right thing even though they evidently must be understanding what they are doing without reference to the complete Newtonian theory. The antikick is such a notion. Rooted in action and counter-action reversibility, rather than $F = ma$; nonetheless, it serves as an important observation which replaces the failing Aristotelian strategy with one which is at least consistent with $F = ma$. This class of prior knowledge may be less obvious than that which manifests itself as bugs, but it is equally important, marking progress in understanding but progress which must not be entirely identified with understanding Newton's laws in a particular context.

The second dimension of refinement which this analysis offers is in introducing context as an important factor in understanding an idea. For example, vector addition of kick to kick is semantically very different from vector addition to kick to established velocity. Except in a view where an abstraction like vector addition is prior to recognition of context-specific semantics, which we should not expect generally to be the case, the differentiation of context would seem most appropriate.

More than just including different contexts, our genetic task analysis suggests natural patterns of development in which context specifics might be peeled away from the abstracted ideas. The transition from Early and Late to Compromise (both involving versions of vector addition) seems little more than leaving the specific naive dynamic feedback rationalization behind in favor of a purer Newtonian conceptualization of combining motions. Presumably, this becomes possible as the notion of velocity, as an entity becomes more clearly developed and hence available to participate in explanations. A second example of abstracting from limited contexts is the transition from Antikick to Generalized Antikick wherein canceling is generalized from pure start-stop situations to ones where possibly interfering initial velocities and intervening kicks can be discounted.

5. NAIVE PHYSICS AT THE UNIVERSITY

There is good reason to believe that naive strategies, such as those indicated in the learning paths chart, are not only relevant to elementary school students but to any physics-naive subjects, and even physics novices. Viennot (1979) has shown that a tendency to identify force with velocity—very much like our "Aristotelian expectation" that things go in the direction they are pushed—is widespread. It occurs in roughly 50% of students, even to the third year of university. Trowbridge (1979) has shown a remarkable inability in university students to deal with the concept of velocity which might parallel the late appearance of velocity-centered strategies.

Previous work, such as cited above, has not attempted to track the possible evolution of knowledge, which is the intent of our genetic task analysis. Even if physics-naive adults or even novice physics students suffer misconceptions like the elementary school students, one may be legitimately skeptical about whether any implications about successful paths to expert knowledge may be drawn from the learning paths chart. Perhaps, one should assume that such naive knowledge simply fades away and is replaced by "proper" Newtonian understanding in the course of schooling. This section suggests that ignoring the naive knowledge state is at best problematic, and at worst tantamount to ignoring important parts of the learning which produces effective expert encoding of the physics. It is based on a protocol of an M.I.T. freshman who, after a year of high school physics and essentially all the Newtonian mechanics in the freshman curriculum, still encounters many of the learning paths chart topics in trying to understand and control the dynaturtle. This is a bright student who knows the formalism (e.g., vector sum) and simply does not think to, or know how to apply the formalism in this phenomenological context. More than that, near the end of the protocol, she appears to be learning the relationships between her continuing "naive" perceptions and her formal knowledge which allow her to see how the phenomena she perceives relate to Newtonian mechanics and how Newtonian mechanics corrects or improves her phenomenon-based insights. After discussing the protocol, we give an interpretation which suggests an important or necessary role for the "connective" knowledge between naive and school physics. Readers interested in more detail than that contained in the synopsis may check the appendix which gives substantial quotations and pictorial representations of the games played.

5.1 Synopsis of the Protocol

The subject (here called Jane) began the task by receiving an explanation of the dynaturtle and Target game in terms of it being a frictionless situation controlled with aiming and kicks or hits. She had been interviewed on five previous occasions, and it was clear that the subject matter was physics. For about half of the one hour, 24-game protocol, Jane used strategies from the learning paths chart almost exclusively. Indeed, her work was dominated by an insistent and explicit use of the Aristotelian Corner strategy. She never used any velocity-centered strategies until after a watershed observation that Aristotle was failing. That critical observation was accompanied by the remark that, after all, it made sense that an object wouldn't lose all its pre-kick motion just because you applied a kick. Subsequently, Jane used two velocity-centered strategies, compromise and generalized antikick, although the ordinary antikick (which can have a velocity-centered interpretation) and early/late persisted into the second half of the protocol. Her use of previously described topics can be summarized as follows:

1. Early verbalization of antikick (before Game 1) and later use (games 19–21). This was only used in the Corner situation to effect a Newtonian Corner, and it was used in place of simple vector addition (diagonal backward kick) even though by that time she showed she was capable of understanding the straightforward vector addition.
2. Abundantly clear and very robust use of Aristotle (games 1 and 5–10).
3. Probable use of Trajectory (games 2 and 3).
4. Certain use of Early/Late conceptualization (games 2, 3 and 15).
5. Probable use of Late Implies Harder (Game 3).
6. Certain use of Compromise (games 11, 12 and 15).
7. Probable use of Generalized Antikick corrupted by giving inverse kicks directed away from the target rather than away from the direction of kick to be canceled (games 13 and 14).

Jane was prompted several times to try to explain what was happening in terms of physics. Though she mentioned such relevant concepts as vector addition and conservation of momentum, she could not explicitly make the connection. For example, soon after her watershed observation of Aristotle's failure, it was suggested she might try drawing a vector diagram, but nothing came of it. Neither did she mention conservation of momentum to justify her post-Aristotelian insight that it was "reasonable" for some pre-kick motion to last through a kick.

Finally, very late in the protocol, Jane made the connection between vector addition and the phenomenology of the dynaturtle. What we consider a crucial stage followed immediately; she proceeded to reflect on and reinterpret her intuitive strategies in terms of the vector formalism. For example, she saw that compromise will usually work as a qualitative version of vector addition. Further, she continued to use her intuitive strategies (the ones which succeeded) through the entire interview. Her last self-proposed strategy involved antikicking to turn a corner, even though the formally posed problem, "What kick (vector) should I give to turn a corner?" yields a diagonally backward solution.

5.2 An Interpretation

Jane's protocol is striking from two related but distinct points of view. We have already discussed the remarkable similarity of her cluster of strategies to those exhibited by 11- and 12-year-old children. But what is equally remarkable is the fact that she did not, indeed for a time could not, relate the task to all the classroom physics she had had. It is not that she could not make the classroom analyses; her vector addition was, by itself, faultless. It is more that her naive physics and classroom physics stood side by side but unrelated, and in this instance, she exercised her naive physics. How are we to make sense of that?

One possibility is to assume that the two types of physics operate within two significantly different representational schemes. One might imagine classroom physics operating within a conscious symbolic scheme typified by discrete entities with well-defined and explicit relations, whereas naive physics might operate in a less integrated way—more like Piagetian action schemes: some gross situational feature cues reactions and expectations, often with no conscious trace or even, for any practical purpose, any discernible internal structure.

Such a view is too speculative to pursue now. Furthermore, we think it advisable not to begin analysis by treating naive and expert knowledge in essentially different ways; that ensures the learning process will appear mysterious. What follows is an attempt at interpreting the relationship between Jane's intuitive and classroom physics which makes no *ab initio* assumptions distinguishing the two types of knowledge. It is intended to be minimal in that it does not propose any specific cognitive mechanisms.

The intent of a genetic task analysis is to identify components of preexisting knowledge which can or do become involved in understanding something like Newton's laws. Our expectation is that conventional methods of assessing what is the knowledge content of subject matter, what we have characterized as abstract task analysis, produce an illusion in terms of the organization and newness of the knowledge involved. Thus, one speaks of "the concept of velocity" as if it were a new entity based perhaps on the logical components of its definition; whereas, in order for that concept to function, as is evident in Jane's case, one must have at least some way of interpreting the naive phenomenology of motion as it relates to that "pure concept." Another way of saying this is that even if one has a parsimonious description of a concept, the way that that knowledge is functionally *encoded* will involve a confluence and complex orchestration of a large number of partial understandings, with many of them based on previous knowledge. We will use the term *distributed encoding* to emphasize this fragmented view of knowing which we take as important to following the genetic lines of understanding.

To elaborate the idea of distributed encoding, we can consider a few commonsense classes of knowledge which contribute to functional understanding, yet are sometimes regarded as ancillary. (1) With regard to experts, no one doubts that they know many special cases which, while those results would follow from the general theory, are used so often that they become automatic and separately encoded. Deduction from the general theory becomes irrelevant for most purposes. One might call this example of distributed encoding pragmatic encoding to indicate that while redundant from a logical point of view, the special case is important in practice to fluid expert behavior. (2) A subtler version of distributed encoding would be techniques particular to some context which allow a general idea to be applied there. Though two problems may be isomorphic, seeing or constructing the isomorphism might require additional knowledge. When an expert says a novice does not really understand an idea, it might well mean that the expert knows a context in which the student will not be able to apply the idea, i.e. the student lacks a way of interpreting the context so as to see the relevance of the idea. (3) Another example of distributed encoding could be a qualitative version of the general theory which might specify only the kind of result one might expect from a detailed application of the theory. This would be important in planning-like activities such as deciding whether and how to apply the theory. (4) Finally, examples and counter-examples, say of a mathematical theorem, play no role in formal mathematics; yet, they play a crucial role in remembering, and the functional understanding of experts and novices alike (Michener, 1978).

With this fragmented view of knowledge, control is an extremely important issue. Speaking loosely, the knowledge system must know when to apply which bit of knowledge. If we take the societal view of mind espoused by Minsky and Papert (Minsky, 1977), and more recently Lawler (1981); and consider these bits of knowledge as independent agents, then each must know when to *defer* to which one, in order to complete a "computation." Although we will not assume control is distributed in this way (what we say is consistent with a centralized executive control), we shall acknowledge the agent metaphor by calling the control knowledge of when to use what ideas *deferences* or *deference links*.

In terms of the classes of distributed encoding mentioned above, one can identify various stereotypical kinds of deference. Examples are known to be examples of some thing; they are understood to exhibit some characteristics of the general case, but are not taken to be definitive. Pragmatic encodings must defer to more general considerations for justification. Qualitative versions are associated in appropriate ways with quantitative refinements (and vice versa).

In Jane's case, when one looks for distributed encodings of vector addition, one finds an interesting state of affairs. Although there appears to be a great deal of what one might be tempted to call distributed encoding (antikick is a special case which Jane uses even after she knows how to replace it by the more general process of vector addition; Aristotelian Corner is a first rate counter-example[7]), these are precisely what she is learning about dynaturtle, and not what she knows about vector addition and Newton's laws. Furthermore, deference is not automatic. In particular, none of those understandings initially defer to "head-to-tail" vector addition formalism, or any of the other classroom physics.

In terms of the general notions of distributed encoding and deference, we understand Jane's case as follows:

1. She knows vector formalism (head-to-tail addition, components, etc.) and F = ma on a certain, mostly symbolic level. This is the "general theory" in the above discussion.[8]
2. She does not have distributed encodings of physics relevant to the gross phenomenology of motion exhibited in the target game that are sufficient to allow her to think of the game in terms of her classroom physics.
3. Instead, she has the means to marshal and develop from her naive physics in relatively short order a rather elaborate collection of understandings which could occupy the place of appropriate distributed encodings of her classroom physics. For these purposes, "naive physics" must include commonsense reasoning and general ideas about causality, canceling, conservation, etc. which might not, except in a genetic sense, be thought of as physics.
4. Thereafter, what remains is establishing appropriate deference links, which is what she appears to be doing in later stages of her protocol (most noticeably in game 18), explaining Compromise, etc., in terms of vector addition.

Point 3 is particularly interesting from a pedagogical point of view; it is meant to point to the learning going on as a student plays with a dynaturtle, independent of formal training. It is, in fact, a possible stage of knowing Newtonian mechanics discussed in Section 4.1, where students can deal effectively with crucial Newtonian phenomena, but do not know any of the formalism. One might liken this to Piagetian concrete operational understanding.

---

[7]Somewhat in the sense of Winston's work on learning (Winston, 1973), one might think of the failure of Aristotle as prompting the recognition of a critical difference which helps transform a preliminary description (based on naive notions) into Newtonian physics. One might guess that adding the property "is-moving" to the set of features relevant to predicting the effect of a kick is key learning from this counter-example.

[8]The role of imputed to "the general theory" or "pure concept" here is an artifact of exposition. We can fix this and liberate the notion of distributed encoding from the image of an expert with a dominant general theory, with "frill-like" appendages (distributed encodings) attached. A more symmetric view is one in which distributed encodings all defer to other distributed encodings. For example, head-to-tail vector addition would play a less central role and assume one more parallel to other encodings. The general theory should then more properly be taken to be precisely the collection of distributed encodings, and their mutual deferences. This places distributed encodings such as those dealing with the gross phenomenology of moving objects on a par (with respect to understanding Newton's laws) with symbolic concoctions like "F = ma," as well as we think they should be.

This interpretation provides a relatively clear role for a certain kind of learning (at the level of distributed encodings) which, if we had a good knowledge map of human understanding, would almost certainly be localized closer to common sense and everyday "intuitive" manipulation of the world than might be expected of learning a new, technical field, such as Newtonian mechanics. In Jane's case, we see someone who evidently has not learned much at this level. Her naive physics seem initially to block, rather than facilitate the use of what she has learned in interpreting the direct phenomenology of motion and interaction. She perceives and interprets the world in a way uniformed by, and incompatible with, Newton. A physicist might say she does not understand physics, even though she can add vectors and recite $F = ma$.

Our genetic task analysis and interpretation of the protocol suggest that learning at this intuitive level could well serve as a platform which would facilitate learning at other levels. Furthermore, that could be done in such a way that deference would be straightforward, if not automatic.

In suggesting this interpretation, there are some caveats. We should be careful not to identify experiential learning with learning at this intuitive level. In the first place, it seems few subjects, if any, had learned much characteristically Newtonian from dealing with the everyday world. Moreover, this work has suggested interventions, such as proposing the Combining Kicks thought experiment, which may be as useful or even more useful than "playing around."

More importantly, we must be aware that the mechanisms by which students usually come to expert status are probably not so clean as our interpretation of Jane's dynaturtle experience might suggest. Though it seems unreasonable not to assume that common sense and experience with moving things around is an important pool of knowledge which contributes to an expert's understanding of physics, especially in terms of distributed encoding, we have no way of knowing how much the set of dynaturtle topics reflects the general case. Neither do we have a way of assessing the contribution of annotation of naive strategies (as in point 4 above) to establishing appropriate deference.

5.3 Remarks on Future Work

All of the above caveats are versions of the hypothetical elements entailed by the concept of a genetic task analysis first mentioned in the introduction, namely that it is difficult to draw definite conclusions about which pools of knowledge might or must have genetic lines into expert understanding. While such hypotheticals may be taken in the context of cognitive studies to constitute obvious future research problems (e.g., is antikick part of the distributed encoding of physicists trained by conventional means?), from the perspective which originally motivated this work (studying learning with the intent of improving instruction), they are less obviously the next tasks to be taken on. One might instead take the conceptual difficulties documented in student's understanding Newtonian mechanics and dynaturtle's apparent engaging of some of those difficulties at face value, and see the next step as one of educational engineering. Though the genetic task analysis proposed in Section 4 may not be *the* route, it seems to be a viable route, and one might concentrate on broadening it and finding ways of ensuring its reliability. Still, there are tactical decisions. One may choose to incorporate insights from our genetic task analysis into a reformulation of the way one explains Newtonian mechanics, for example, incorporating genetically germain, but logically redundant concepts like antikick (diSessa, 1980). One may choose to add an automated tutor to the target game in the spirit of Goldstein (1977), and Burton and Brown (1979) which could use a genetic task analysis as part of its student model. Finally, one may choose to fine tune the dynaturtle environment (the target game was, after all, only a first guess at a productive environment) to be more effective, reliable and better linked to standard teaching. White (1981) has taken this approach.

6. CONCLUSION

A genetic task analysis is intended to be a fundamentally different slicing of a domain than that achieved with an abstract task or conceptual analysis; genetically antecedent, partial understandings replace logical prerequisites as elements. Not only is this a pedagogically advantageous point of view, but we have argued that it may in fact reflect essential aspects of expert encoding; what might be taken as ancillary or redundant encoding of a concept may serve important functions such as facilitating quick access, providing for robust remembering and allowing higher level, qualitative analysis to aid in planning, etc. It might even be important in allowing the application of the concept to its intended application domain—in this case, the phenomenology of interaction and motion. We have elaborated a view of the knowledge involved in some aspects of Newtonian Mechanics in such a way as to find a natural place for experience such a playing with a dynaturtle, or more precisely, for the interpretations of that experience drawn from naive physics.

A final question remains. Why can't the students' experience in the real world serve the same purpose as an experience with dynaturtle? Why do students come to dynaturtle with deep Aristotelian misconceptions? In part, this may be accounted for in the striking ability of humans to hold theories of their own action which contradict what they do in fact. Dynaturtle's advance over naive experience, then, lies in the explicit and unambiguous actions taken to control it. Experience with dynaturtle is mediated by a very narrow channel of kick and turn commands, as opposed to interpretation of complex muscle actions actually used by humans in moving things around.

But a better explanation as to why the real world doesn't teach Newtonian mechanics probably lies in understanding how good a non-Newtonian theory like "kicking in the direction of intended motion" can be. This certainly suffices for cueing up a billiard ball and works whenever impulse dominates existing momentum. Further, in many circumstances, one simply *arranges for* the theory to work. Compare a soccer player who stops a ball as a matter of course before kicking again to the Newtonian Corner strategy.

Finally, in the real world friction has two confounding effects: one supporting Aristotle and one denying Newton. By rapidly bringing velocities near zero, it allows an Aristotelian plan to be more generally effective, thus mitigating the need for refinements. More fundamentally, friction denies Newton's first law by its very presence; the world is *prima facie* non-Newtonian. Since friction is omnipresent and with no visible agent causing it, why should one either implicitly or explicitly treat the "dying away" of motion, so much like other inescapable things, as other than a primitive phenomenon (law) of nature? It is only by coming to understand the Newtonian stance that one even acquires a reason to separate friction as another force to be included in the analysis. And beyond the first law, of course, the second law doesn't work without frictional forces being explicitly included. Summarizing this line of reasoning, a Newtonian frame of analysis seems necessary to make sense of the notion of friction as a force, rather than as a fundamental and universal phenomenon intrinsic to motion. Yet, a Newtonian frame is only possible after one has separated out friction as a force to be added to the analysis. Galileo and Newton's escape from the bind truly betrays their genius. In the present case, the bind is not inescapable, as we can simply remove the confounding element from the (simulated) world. Dynaturtle is a pure representation of Newton's laws, unfettered by friction.

In closing, I would like to return to the educational motivations of this work and point to the Aristotelian-Newtonian controversy as a "play within a play" framing the context and importance of this kind of research. As with our naive subjects' assumption about dynaturtle, most teaching seems to assume no dynamical state on the part of the students, that there is not much of interest in their present knowledge for predicting future learning trajectory. "Pushing in the direction you want to go" seems to work pretty well, especially if one pushes hard enough. But what this kind of work begins to reveal to us, as dynaturtle revealed the notion of momentum to our subjects, is that there is a rich and complex knowledge state that one can use to good advantage in attaining pedagogical aims. It may in fact be true that in certain semantically rich domains, a student's initial state dominates the perturbations we can apply as teachers. The depth of our understanding of the student's knowledge state and our cleverness in engaging its subtleties may then determine the ultimate success or failure of our teaching efforts.

APPENDIX

Chart 2 represents the first 20 games in a 24-game, one hour sequence in an audiotaped protocol of a freshman MIT student introduced to the dynaturtle Target game. The computer was programmed to record each keystroke and timing thereof to allow playback and other processing of that part of the protocol. While both kicks (arrows) and trajectory are shown here, neither was displayed for the student.

Our mode of analysis will be simply to parse the protocol (roughly at the grain of one game) into episodes which we interpret as being dominated by one of the learning paths chart topics. In addition, we have selected, annotated and interpreted pieces of text from the protocol mainly on the basis of their relevance to determining to which topic (if any) the game belonged. No attempt has been made to verify any of the proposed topic-to-topic links from this one case. Wedged brackets are intended to mark and explain omissions from the text. Square brackets are intended to mark more interpretive commentary.

1. ANTIKICK

Start of Protocol

Interviewer: *This is going to be a little creature that you drive around by giving commands. All you have to do is hit a key and it does things.*

Jane: *Forward, backward.*

I: *Yeah, except the commands aren't forward and backward. Have you seen turtles?*

J: *No, but I've seen the spaceship ones* [video games]. < *She explains further about playing gallery games.* >

I: *I see. This is almost like that.* < *Demonstrates R and L.* > *And the other thing you can do is hit a K which stands for "give me a kick in the direction I'm facing." He behaves as if there isn't any friction so that if you give a kick, he's just going to keep going that way.*

J: *So that you've got to turn him around and kick him back to stop.* [A clear enunciation of antikick.]

I: *Or whatever.* [The interviewer does not wish to be committal about the effectiveness of antikick and proceeds to explain the object of the game: to hit the target but with as small a speed as possible.]

2. ARISTOTELIAN CORNER STRATEGY

[Jane's first game started with a kick aimed as closely as possible at the target and then proceeded with a series of kicks aimed directly into the target, i.e. a series of (non-right angle) Aristotelian Corners (see Game 1). At this stage, she does not explain what she's doing and evidently does not see that Aristotle is failing.]

3. TRAJECTORY STRATEGY— EARLY STRATEGY AND LATE IMPLIES HARDER

Game 2

<Plays game>
I: *Tell me what you're doing.*
J: *It's just, <pause> more the way I feel. . . . The first time I had started turning it back too late. I let it go forward before I started changing its direction, kicking it, making it curve.* [She seems to have switched to Trajectory between games. Note the use of "curve" language and kick aimed clearly away from the target.] *So this time I started changing it back a little sooner.* ["Sooner" is a key clue to use of Early Strategy. Apparently, she reinterpreted the results of Game 1 within the Trajectory perspective and applied the Early conceptualization and feedback loop to imply the corrective action "do the same, only earlier."]
I: *I see.*
J: *I guessed exactly when because I wasn't sure yet. . . .*
I: *So you were just estimating when to*
J: *turn it. Cause the first time I'd done it too late so I tried a little earlier this time.* . . . [This is a clearer pronouncement of Early.]

Game 3

J: *That time I over-corrected,* [meaning corrected more] *because I knew I wasn't in as good a position so I had to start bringing it back real fast. <She illustrates with her hands the path not turning enough and a kick to turn it more.>* [She is using Late Implies Harder to suggest a second turn-kick when the first evidently didn't do the job. Note the extra turn-kick compared to Game 2. Also note rhythmic (equally spaced) application of turn-kicks, again characteristic of Trajectory, in both games 3 and 4. Note further how clearly these attempts would seem to refute Aristotle if she had it at all in her mind to do that! The second kick in game two is particularly clear. One reason she does not see how clearly Aristotle is being confounded is that she has switched strategies to Trajectory and is therefore focusing on different features—amount of turning of path rather than heading of the dynaturtle.]

4. ARISTOTELIAN CORNER STRATEGY

[The next five games offer a remarkable sequence of attempts to make Aristotle work. Her discussion leaves no doubt as to the strategy used. By Game 7, Jane is aware of a problem more clearly, but, somewhat frustrated, attributes it to "mis-aiming." Finally, at Game 10, she offers a hesitant question if dynaturtle might actually *not* go in the direction it's pushed.]

Game 5

I: *Tell me how you're going to do better.*
J: *Well, first of all I'm going to avoid giving it more kicks, if at all possible. So I'm going to try to keep it going along a 45° angle the whole way. And that way I won't have to give it another kick.* [These three sentences relate to a topic not listed in the learning paths chart which is essentially that speed equals number of kicks.] *When I turn it I have to give it another kick to get it to go in that direction.* ["A kick to get it to go that way" is classic Aristotle.] ... *If I can keep it going straight at the target I should be able to avoid giving it more velocity, or whatever, to get it to go.* [Here she is reverting to the topic of the first three sentences above which involves equating speed with number of kicks.]

Game 6

J: *That time, ok, I'd given it the one turn, one push to the right.* [This is the aiming for startup kick.] *I let it go a way until I estimated that if I gave it another turn to the right it would be headed straight toward the target. Then I gave it one kick to get it going in that direction hoping it would hit.* ["Aim toward target, then kick to get it going that way" is pure Aristotle.] *It did. If it had been obviously off I would have tried to adjust it.*

Game 8

J: *No! Darn it, I'm misjudging all these.* [She is attributing the failure of Aristotle to misaiming.] <*Starts Game 9*> ... *I'm starting two over* [two right turns over from straight up] *and then I'm letting it go for a while. And then I'm bringing it back one* [one turn left] *and waiting until I think it's pointing right in the circle.* [A refined Aristotle, "wait until it's pointing toward target" instead of actively turning toward target, but still clear Aristotle.]

Game 11

J: *Can I ask, like, ok, if I, it seems like it doesn't head straight in the direction it was pointing when I gave it the second kick.* [Finally, she sees Aristotle fails.] <*She explains setting up for a final kick.*> *But when I did give it a kick, it seems like there was a little bit of this* [indicating direction of motion just before kick] *left.* ["A little bit of initial motion left after kick" seems a proto-momentum conservation.] *It wasn't going exactly there* <*indicating kick direction*>, *but kind of in between.* [The topic has shifted to Compromise. Note that the proto-momentum conservation is associated with the first use of a velocity-centered strategy.]
I: *It's supposed to be really simulating a tennis ball or something going along, except without friction, so that the kicks that you give it are just like what would*

happen if you gave little kicks to a tennis ball. You should really think that it's supposed to be physics. [The interviewer constantly prompted Jane to use her classroom physics.]

5. COMPROMISE

Game 12

J: *It isn't like it's starting from rest.* [This is a perfect characterization of Aristotle's failure.] *Which means I'm going to have to make more corrections. ... Which means I'm going to aim a little bit lower than the target.* [''Lower'' is opposed to ''above the target'' which is the direction of motion of the dynaturtle.] *If I aim a little bit lower than the target, it should end up right at the target.* [This is clear Compromise, ''above target'' (pre-kick direction of motion) + ''lower than target'' (kick direction) = right on target.] ... *It works if I aim just a little bit below the target.*

I: *It seems to take a compromise path or something like that?* [Interviewer is testing her use of Compromise.]

J: *It seems that way. I'm trying to think of any equations or anything that explain that. Except that it seems, like, logical that it* [dynaturtle] *would have some of it* [pre-kick motion] *left over. ... I mean you can't just give it a kick and expect it to ignore that it was already going.* [Again she's exhibiting a focus on continued existence of previous motion through a kick.]

I: *Can you think of any equations, or* [Interviewer encourages her to think physics, sensing she is groping in that direction.]

J: *The conservation of momentum,* [Has she got it?] *except, well that was just in my head from this test I had this morning.* [No.] *Maybe using vectors rather than just magnitudes. ... but even that isn't a very good explanation.* [Of course, both momentum and vector ideas are relevant. She seems to have that sense, but cannot make an explicit connection.]

I: *That doesn't sound very good?* [Again encouraging physics thinking.]

J: *Not really. <long pause> No. It just* seems, *when you hit something when it's going in one direction, it's not going to completely change its direction. But I'm lost for an explanation at this time.*

I: *That vector stuff doesn't help?* [One more try at eliciting physics.]

J: *It should. <Starts drawing. No Clarification arises.>* [She seems incapable of making the connection on her own.]

6. GENERALIZED ANTIKICK

Game 13

I: *Do you have any theories about getting to hit the target slower?*

J: *<She asks if the turtle can enter the target while it's facing backward.>* [This may be the same kind of concern with Aiming Independent of Motion that more subjects demonstrated in the Antikick context.]

I: *Yeah, it can enter the thing no matter which direction its facing. Entering it just depends on the velocity it has.*

J: *OK, then I have an idea. ... I'm going to give it two kicks, then, right before* [before hitting the target] *I'm going to turn it around backwards and kick it in the opposite direction to cancel one of them.* [She wants to cancel one kick leaving another; the topic is Generalized Antikick. Note explicit reference to canceling] *Then it should enter going only one. <Plays game.>* [It turns out that "opposite direction" gets implemented in the peculiarly Aristotelian fashion, "opposite the direction toward the target!"] *That one didn't work at all.*

Game 14

<Plays game 14.> [She starts again with canceling in mind (first four kicks are the same as Game 13), but then reverts to kicks toward the target.]

J: *<She facially evidences surprise at her good score> ... Obviously the sum of the vectors all added up to one* [Her score was one kick's worth of velocity. It sounds as if she now understands what's going on—"vector sum."] ... *I was just watching the thing and hitting it as I felt I needed it ...* [But she explains that she was following an intuitive strategy in actually playing the game.]

7. COMPROMISE

Game 15

J: *I have an idea now. When I give it a third kick that was supposed to be in the opposite direction, it headed more straight upwards. This time I'm going to wait a little longer to give it a kick,* [This is a fleeting use of Early] *<brief pause> aim it a little bit below the target and hopefully the direction will be right.* [Early rapidly gives way to the better conceptualization, Compromise. To understand her use of Compromise, consider the velocity after the second kick in game 14, bearing 45°. To this one adds a kick at bearing −120° and gets a result between the two, about 30°. (Presumably she accepts that the velocity direction, for some reason, gets a high weighting factor in the compromise.) So she expects to substitute a kick at only −90°, and gets a result even closer to the original velocity direction, 45°.] *<Plays 15.>* [This fails and even gives the opposite result. (Figure 6 shows the vector addition.) She follows the failure with emergency kicks toward the target.]

Game 16

Game 16 is roughly the same as Game 15.

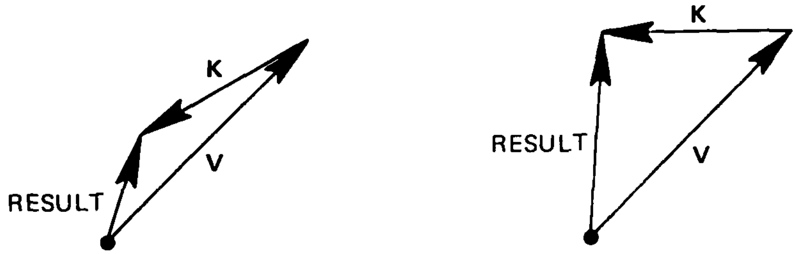

Figure 6. In contrast to the Compromise expectation, rotating a kick clockwise can cause the resultant velocity to shift in a counter-clockwise direction.

Game 17

<*Plays 17.*> [Game 17 appears mostly playing with Aristotle and anti-Aristotle (kicks directly away from the target).]

J: *Maybe I should just play this game by goofing around. This isn't fair. How could I get such a low score, just by goofing around.* [She admits to playing around and is somewhat upset at getting a good score by doing so rather than by applying a consistent strategy.]

8. NEW TOPIC—VECTOR ADDITION!

Game 18

[For the first time, the topic of vector addition, previously mentioned several times, has center stage. What is particularly interesting is her attempt to interpret intuitive strategies in terms of it.]

I: *Could you explain what's happening?*

J: *I think it's adding up the vectors of the velocities. . . . It adds them together to find the final velocity. Then the direction will be different. I think it will account for the change in the, not the exact direction of the second, but a compromise between the two.* [She sees how vector addition means Compromise usually works.] *And also I think maybe that the second, like the impulse would give it twice the velocity, but adding them up maybe it doesn't go . . . maybe the magnitude won't be two times.* [She also has a glimpse of why, because of vector addition, her earlier theory of speed = number of kicks doesn't work.] <*She draws and gives a good explanation of adding impulse velocity to established velocity.*> *In other words, I'm just adding up all the different velocities in all the different directions, then that's what's going to happen.*

9. ANTIKICK AND NEWTONIAN CORNER STRATEGY

Game 19

I: *Do you think you could make it fly straight up here* <*indicating the height of the target*> *then come in at 90°?* [The task of making a corner is proposed without any suggestion of method.]

J: *<Starts game.> . . . Now I'm going to turn it around and give it a kick in the opposite direction when I think I'm parallel to the target. Now I'm going to turn its direction and kick it. . . . Yeah, that worked out really well.* [This is clear Newtonian Corner. Notice that she did not use her vector understanding to ask which kick, when added to vertical velocity, would give a resultant horizontal velocity. That would have prompted the "diagonal backward kick." Notice further that she had not previously used antikick successfully.]

I: *Could you do that faster?*

J: *I would just give it a bunch of kicks and then turn it around sooner and give it the same number of kicks to stop it. I might misjudge it, but I know how I would do it.* [She is very confident of her reasoning, although she is not drawing vectors.]

Game 20

J: *I gave it five kicks but you could give it any number.*

Games 21–24 involve strong intervention on the part of the interviewer and are of no interest here.

CHART 2

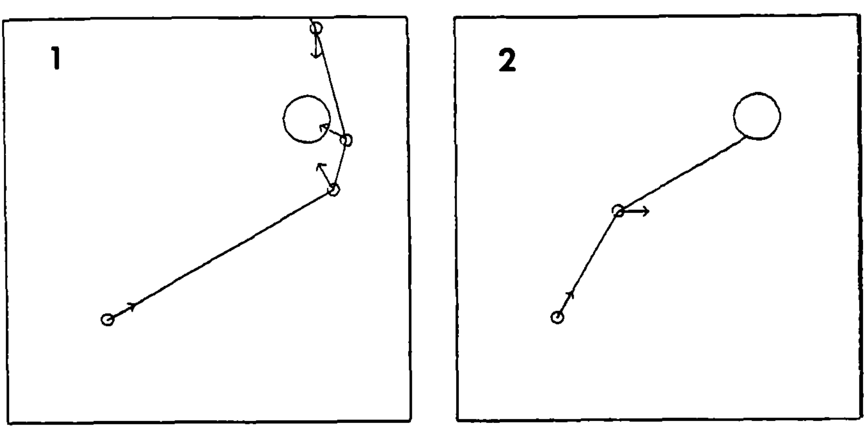

CHART 2 (continued)

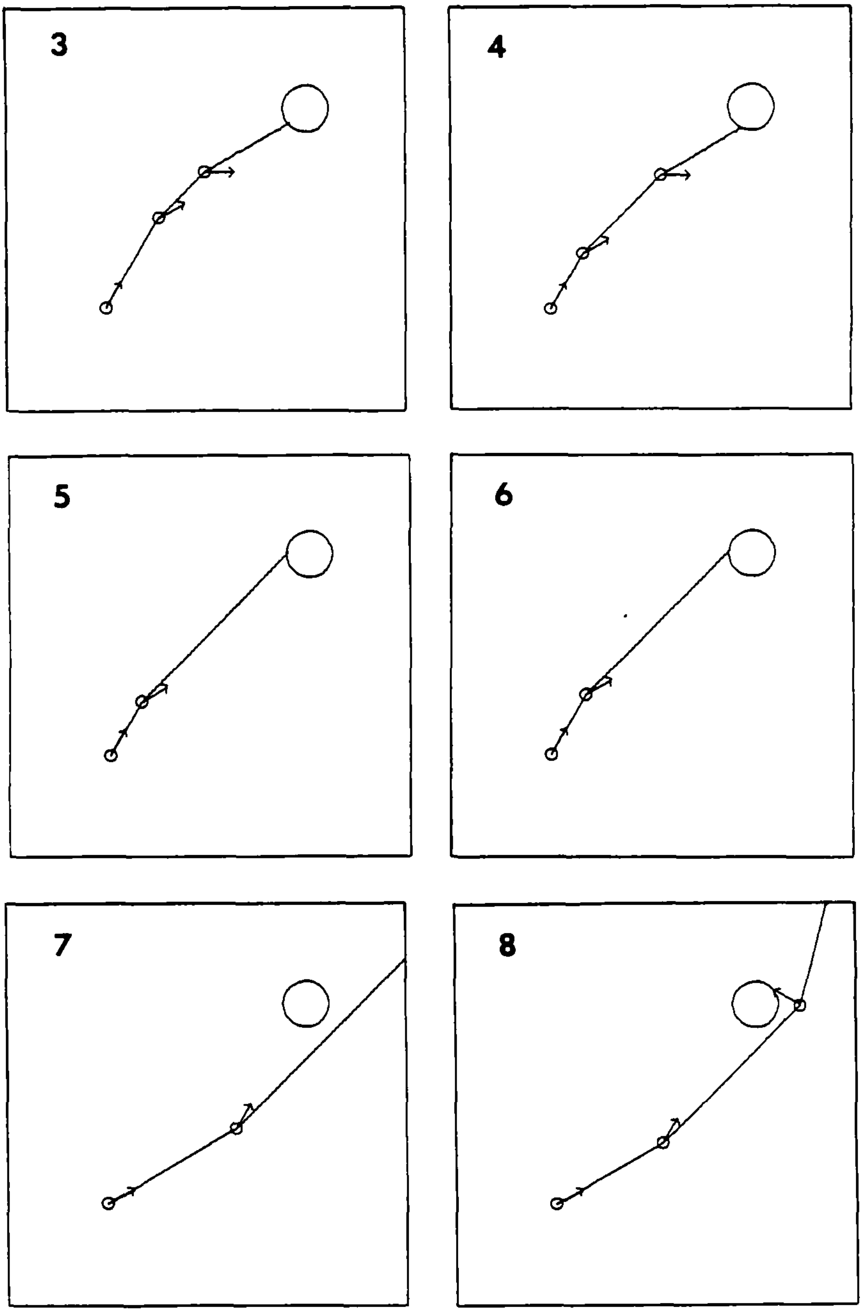

CHART 2 (continued)

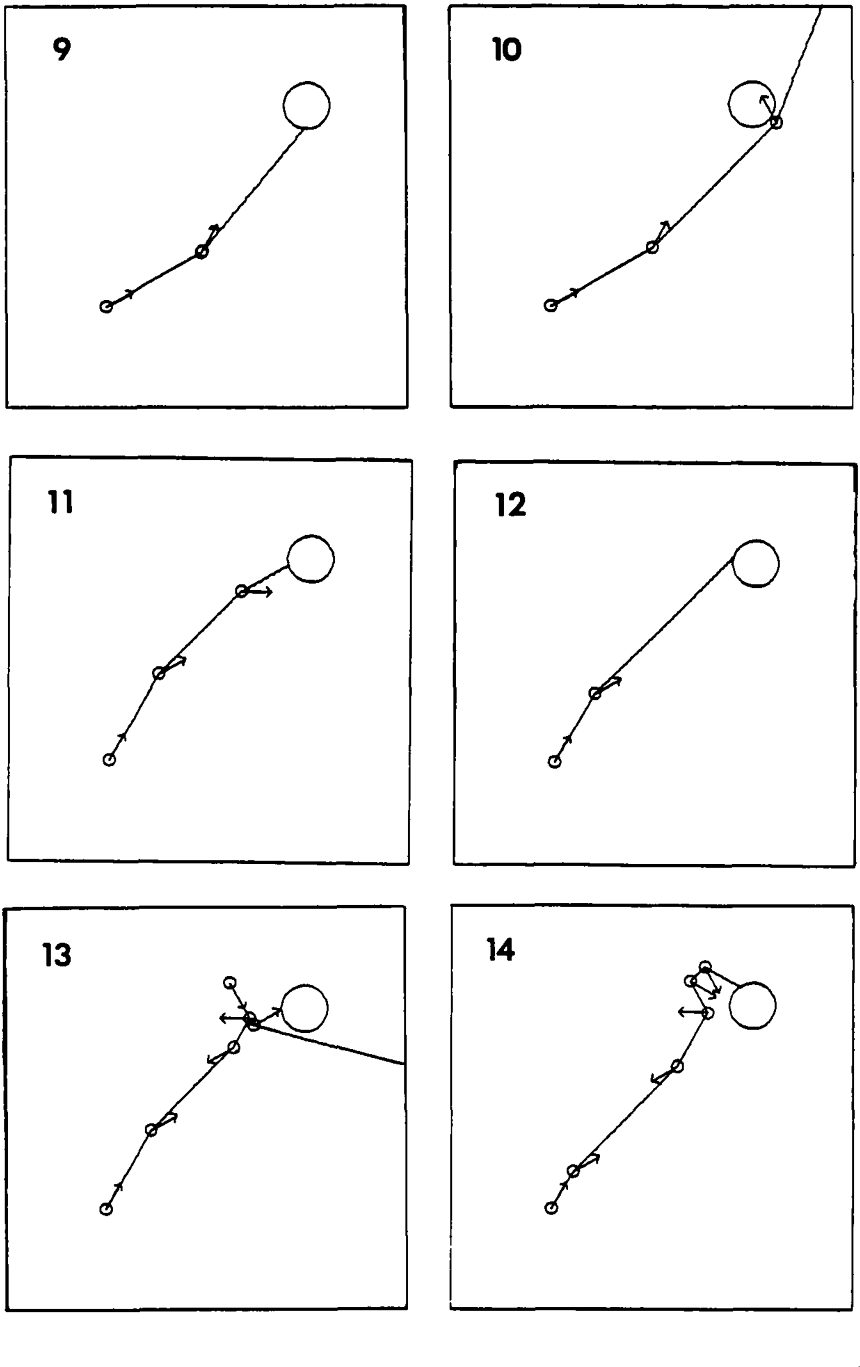

CHART 2 (continued)

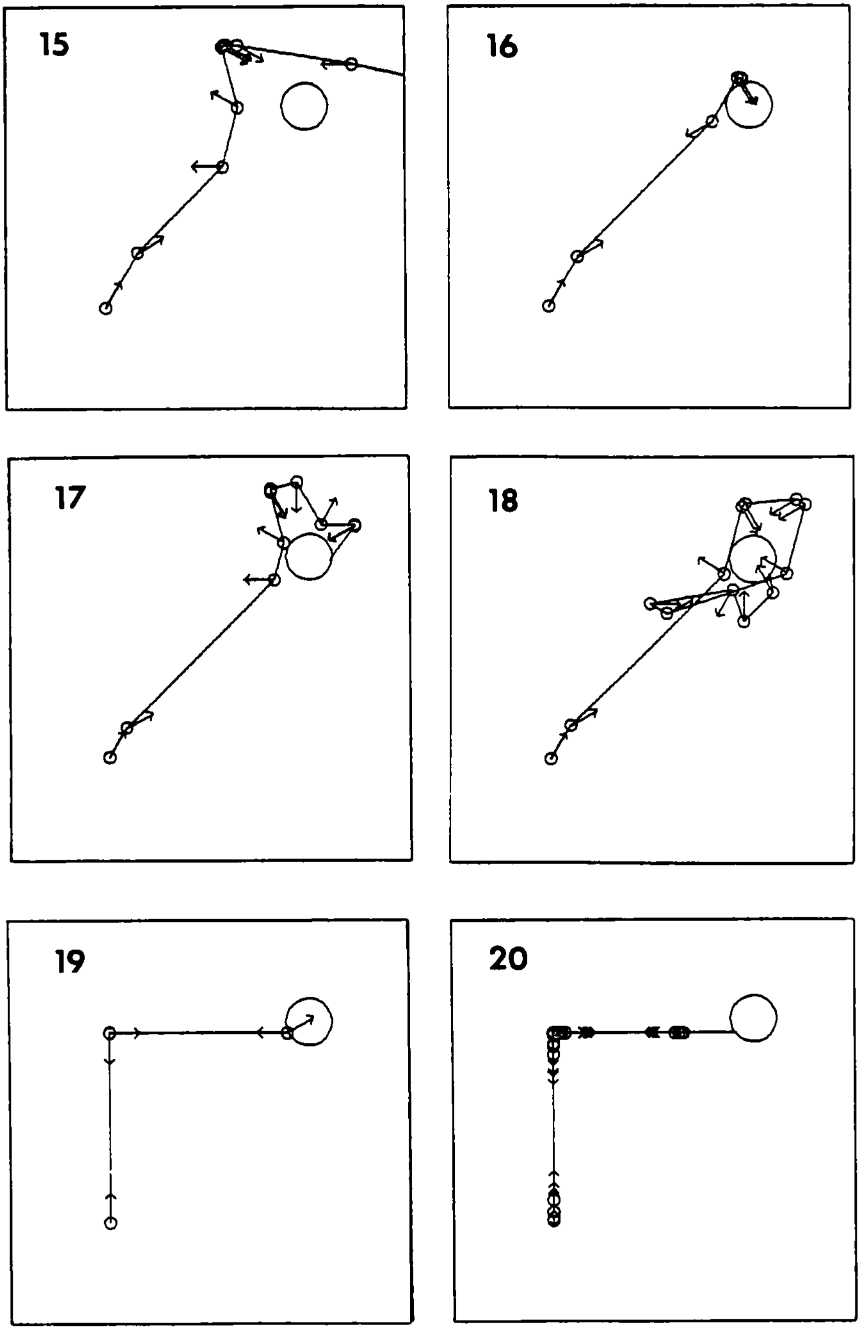

REFERENCES

Burton, R. R., & Brown, J. S. An investigation of computer coaching for informal learning activities. *International Journal of Man-Machine Studies,* January 1979.

Clement, J. Intuitive, non-Newtonian views of motion. Unpublished mimeo. Department of Physics and Astrophysics, University of Massachusetts, Amherst, MA, 1977.

Cohen, H. The art of snaring dragons. (AI Memo 338.) Cambridge, MA: M.I.T. Artificial Intelligence Laboratory, 1974.

diSessa, A. Intuition as knowledge. Unpublished manuscript, Cambridge, MA, 1978.

diSessa, A. Computation as a physical and intellectual environment for learning physics. *Computers and Education,* (Vol. 4, No. 1), 1980.

Goldstein, I. The computer as coach: An athletic paradigm for intellectual education. (AI Memo 339.) Cambridge, MA: M.I.T. Artificial Intelligence Laboratory, 1977.

Lawler, R. The progressive construction of mind. *Cognitive Science,* 1981, *5,* 1-30.

Lawson, R., Trowbridge, D., & McDermott, L. Students conceptions of dynamics. Unpublished mimeo. Seattle, WA: Department of Physics, University of Washington, 1980.

Michener, E. R. Understanding understanding mathematics. *Cognitive Science,* 1978, *1,* 361-383.

Minsky, M. Plain Talk about neurodevelopmental epistemology. *Proceeding of the fifth international joint conference on artificial intelligence,* Boston, MA, 1977.

Papert, S., Watt, D., diSessa, A., & Wier, S. Final report of the Brookline Logo project, (Vol. 2). (AI Memo 545.) Cambridge, MA: MIT Artificial Intelligence Laboratory, 1979.

Shanon, B. Aristotelianism, Newtonianism and the physics of the layman. *Perception,* (Vol. 5), 1976.

Trowbridge, D. *An investigation of understanding of kinematical concepts among introductory physics students.* Unpublished doctoral dissertation in the Department of Physics, University of Washington, Seattle, WA, 1979.

Viennot, Laurence, *La raisonnement spontane en dynamique elementaire.* Paris: Hermann, 1979.

White, B. Y. *Designing computer games to facilitate learning.* Unpublished doctoral dissertation in the Department of Electrical Engineering and the Division for Study and Research in Education, M.I.T., Cambridge, MA, 1981.

Winston, P. H. Learning structrual descriptions from examples. In Winston, P. H. (Ed.), *The psychology of computer vision.* New York: McGraw-Hill, 1973.