diSessa, Andrea A. (1988).Knowledge in Pieces. In Forman, G. and P. Pufall, eds, *Constructivism in the Comuter Age*, New Jersey: Lawrence Erlbaum Publishers.

Scanned for a pset buddy 5 Dec. 13

Chapter 4

Knowledge in Pieces

Andrea A. diSessa
University of California, Berkeley

Nobody thinks clearly, no matter what they pretend. Thinking's a dizzy business, a matter of catching as many of those foggy glimpses as you can and fitting them together the best you can. That's why people hang on so tight to their opinions; because, compared to the haphazard way in which they're arrived at, even the goofiest opinion seems wonderfully clear, sane, and self-evident. And if you let it get away from you, then you've got to dive back into that foggy muddle to wangle yourself out another to take its place.

—Dashiell Hammett

How one intends to use computers to aid learning depends in a dramatic way on what one thinks is important in learning. In this chapter I outline a central theme of my work with computers and learning which follows from certain empirically and theoretically driven predilections concerning the nature of knowledge and its development. The fundamental question is: How do we view the transition from commonsense reasoning about the physical world to scientific understanding? Leaving aside the nonconstructivist "accretion" model—new knowledge by transmission from textbook or teacher—there are still very different views of learning that motivate different approaches to the uses of computers.

My own view is that the transition to scientific understanding involves a major structural change toward systematicity, rather than simply a shift in content. After outlining this view by contrasting it with another that presumes a more evenhanded trade of content from prescientific to scientific apprehension, I will discuss uses of computers that follow more or less directly from the structural-change perspective.

INTRODUCTION

What is the character of the knowledge that people spontaneously acquire about the physical world? How do people think the world operates based on their experience with it? This is a subject on which Piaget and his colleagues have spent many productive years. It has come into focus again in recent years at ages beyond those Piaget usually studied, late high school and early college years. In this setting, there has been less emphasis on cognitive development and more on developing understanding in more formal situations: science and mathematics courses. Intuitively developed physics is revealed in interaction with the concepts and theories that physicists hope to teach.

In very brief summary of this line of work, it seems that intuitive physics is a rather well-developed and exceedingly robust system that can potentially interfere with "proper" textbook understanding. A large set of probes has been developed in which students give relatively uniform, but incorrect or at least nontextbook answers, long into the educational process that is meant to provide a proper understanding of the laws of nature.

To sharpen focus on this phenomenon, I would like to contrast two opposing views of intuitive physics. The first is an example of what I call "theory theories," and holds that it is productive to think of spontaneously acquired knowledge about the physical world as a theory of roughly the same quality, though differing in content from Newtonian or other theories of the mechanical world. Michael McCloskey of Johns Hopkins is probably the most visible of the theory theorists at the present time (McCloskey; 1983; July 1983). He described his research as follows: "We show that . . . people develop on the basis of their everyday experience remarkably well-articulated naive theories of motion. Further, we argue that the assumptions of the naive theories are quite consistent across individuals. In fact, the theories developed by different individuals are best described as different forms of the same basic theory" (McCloskey, 1983, July, p. 299).

McCloskey went so far as to tell us what the core theory is that essentially everyone has: It is a version of the impetus theory developed in the Middle Ages, standing historically between two great landmarks, Aristotle's physics and that developed by Newton. I will provide some details shortly.

On the other side, my own view is that this is a highly misleading representation of the actual state of affairs. Though it gives signs of being quite robust, intuitive physics is nothing much like a theory in the way one uses that word to describe theories in the history of science or professional practice. Instead, intuitive physics is a fragmented collection of ideas, loosely connected and reinforcing, having none of the commitment or systematicity that one attributes to theories.

There are many implications to the dispute, but they become particularly pointed when it comes to educational implications. The garden path and frequently advocated strategy on the theory theory side is to attempt to provoke a theory change: to expose and confront the intuitive theory with evidence and argumentation so that students can switch theories. My own view is that a much broader attack needs to be made. Indeed, "attack" is certainly the wrong word. Not only is a one-by-one attack of the knowledge fragments that constitute intuitive physics a hopeless task, but the only material we have to develop scientific understanding in our students' heads is precisely those same fragments. One must not throw the baby out with the bath water. And, although there is surely some trading of one content for another, the issue of form is equally, if not more important. *Building a new and deeper systematicity* is a superior heuristic to the "confrontation" approach many theory theorists have taken.

In the second major part of this paper I develop images of computer-based pedagogy appropriate to science education in view of the character of intuitive knowledge and its relation to textbook physics. These come in three flavors. The first involves engaging naive knowledge on the level that makes best connection to it: experience. Computers provide an excellent medium for designing activities that build and integrate pieces of knowledge. Integration, however, needs special attention, as we need do more than just let children play with simulations and scientific models of the world.

The second use of computers involves replacing static and abstract formalism of the past with ones that are better linked with intuitive knowledge. Symbols like equations and numbers require substantial internalized knowledge to operate and to connect meaningfully with the "real world." On computers we can craft systems that are at once more expressive of dynamic and interactive aspects of the world, and, because they operate more like real-world systems themselves, are easier and more engaging to learn.

Finally, at a level above both of these described roles, our students should learn more about the nature of the development and integration of knowledge itself so as to better monitor and control their own learning. Computers don't play any single special role here, but instead, as with other educational tools, the goal of developing awareness and skills "at the meta level" will influence in many ways what we should do with technology.

TWO INTERPRETATIONS

I will begin the comparison of the impetus theory theory and my "knowledge in pieces" on some of the former's strongest grounds. This is a context that provides some of the best evidence that there is such a thing as an intuitive impetus theory. In this way, I can give a reasonably compelling, though brief, presentation of the impetus theory and at the same time test my alternative view on more-than-fair grounds.

The Impetus Theory

Consider the following simple problem: What happens when you throw a ball straight up into the air and catch it? A "cleaned up" protocol of a high school or college student might run something like the following: "When you give the ball a toss, you give it a force that propels it into the air. But this force is working against gravity, and as it dies away, gravity begins to take over. The peak of the trajectory is the point at which gravity is just balancing the force you gave the ball, after which gravity overcomes that force and causes the body to fall downward at an increasing rate."

The "force" you give the ball that propels it into the air against gravity is impetus, an internal force that can be imparted to an object, and it is, according to McCloskey, the central actor in the impetus theory that characterizes naive ideas of motion. Impetus has other characteristics. It spontaneously dies away, or it may die away as a result of things like friction. There are some fine points. Notably, there are two kinds of impetus, linear, as in the above example, and circular, which we shall encounter below, but this little sketch highlights in capsule form the main ingredients of the impetus theory theory.

A physicist's analysis of the toss involves only one force, that of gravity acting constantly downward. Any upward force ends when the hand loses contact with the ball. There is no "balance" at the peak of the trajectory, nor any "overcoming" on the way down. There *is* a construct in physics that has some of the properties of impetus. It is momentum. In fact, momentum is transferred to the ball in the initial stages of the toss, and the momentum is, in a way, "responsible" for the ball moving upward. But momentum is not a force, it doesn't die away of its own accord, and it does not combine or conflict with other forces in the way the impetus explanation of the toss suggests. These caveats having been made, impetus does not make a bad preliminary metaphor for momentum, and I will frequently use it in this productive way.

There is no doubt that people sometimes give protocols that look like the above fiction, but the central question here is whether this is indicative of a widespread theory of motion or, instead, might arise in a quite different way.

Knowledge in Pieces

My alternative view (diSessa, 1983) is that intuitive physics consists of a rather large number of fragments rather than one or even any small number of integrated structures one might call "theories." Many of these fragments, which I call "p-prims" (short for *phenomenological primitives*), can be understood as simple abstractions from common experiences that are taken as relatively primitive in the sense that they generally need no explanation; they simply happen. For example, why is it that you get more result when you expend more effort, say, pushing a big rock? There is no ready explanation, nor really any need for one. One has so much experience with things that work like that, that the

phenomenon is encoded simply as an expected event. There is need for furthe thinking only when things *fail* to work in that way.

Let me list some examples of p-prims. Table 4.1 gives a sampling that wil prove very useful to us in returning to the tossed ball.

TABLE 4.1
A List of P-Prims Together with Key Attributes That, in Part, Define Them, and a Prototypical Circumstance from Which the P-Prim Might Be Abstracted and to Which It Applies

Name	*Key Attribute*	*Prototypical Circumstance*
Ohm's Law	Agency (also "resistance")	Pushing a box with variable effort on different surfaces
Force as a mover	Violence	A throw
Continuous force	Steady effort	A car engine propelling a car
Dying away	Fading amplitude	Sound of a struck bell
Dynamic balance	Conflict	Equal and opposite competing forces
Overcoming	"Success"	Greater force overcomes weaker

Ohm's Law is one of the most fundamental and pervasive p-prims. It is really an enlarged version of the "more effort begets more results" primitive mentioned before. It consists of an agent or *impetus* (*impetus* in a different sense than in the "impetus theory") that is exerting some effort to achieve a *result* through some *resistance*. In such circumstances the proscribed behavior is that increased effort begets increased result; increased resistance begets reduced result; and so on. Not only is this the commonsense interpretation of Ohm's Law, which describes the relations between voltage (the impetus), resistance (the resistance), and current (the result), but it also interprets a very broad range of physical, psychological (e.g., "trying") and even interpersonal situations such as "influencing." The key attribute, agency, is one that plays a central early role in intuitive physics, and it has a long and interesting development, though one I cannot describe here.

Force as a mover is a simple abstraction of a throw. It involves a directed impetus, a rapid pattern of effort, then release, and a result in the same direction as the impetus. The result, which is modulated by Ohm's Law with respect to the impetus and the resistance (weight, etc.), may be either a net result (distance) or a more local one (speed).

Continuous force shares a common abstraction with force as a mover in that a directed impetus achieves a geometrically parallel result according to Ohm's Law, except that I believe these two are separately encoded. Such redundancy is typical of intuitive physics and is one of the reasons for its apparent robustness. The two p-prims differ in the pattern of amplitude, which is described as "violence" for force as a mover and "constant effort" for continuous force. Such patterns are an important class of attributes for cuing and encoding of physical p-prims.

That sounds, motion, and so on, all *die away* of their own accord is another phenomenon involving a characteristic pattern of amplitude, in this case, gradual fading. This is also a good case of a p-prim being "relatively primitive" in that, even though it is generally taken to be a fact of life that needs no further examination, people will often find excuses for it, such as competing influences (gravity wears away the linear motion of a rolling ball somewhat like it makes us tired in walking).

Dynamic balance involves a direct conflict between opposing forces. Presuming they are equal, neither gets its way; but if one becomes stronger or the other weaker, the stronger will win out, "overcoming" the other, perhaps with a crescendo of "result." Thus the potential action called *overcoming* is closely connected to dynamic balance as an expected possibility. In general, balance and equilibrium is a rich, salient and very important class of primitive phenomena in intuitive physics. *Being in equilibrium* is frequently given as a primitive explanation for why things are as they are.

Let us take another look at the toss of the ball in terms of these p-prims. The first part of the toss, the action of your hand on the ball, is essentially never described at all because it is entirely unproblematic. The p-prim of force as a mover describes and explains precisely this situation. Descriptions given by subjects are often as interesting for what they don't say as for what they do. Indeed, many of the explanations given by subjects must be expected to be comments on what are seen as problematic or puzzling aspects of a phenomenon rather than reductions to a fundamental set of principles, which is what problem solving means in a more formal context. In this case the existence of force as a mover explains why people never make any analysis of the first fraction of the toss, though it is certainly warranted from a physical point of view.

In contrary manner, the rest of the toss is intuitively problematic. There is evidently a conflict involved in the situation; although gravity wants to cause the ball to go down, it continues upward until the peak of the trajectory. So already p-prims having to do with conflict are cued. Even more, the peak of the trajectory in its commonsense interpretation of "stopped," fairly oozes *balancing* and equilibrium. The down side of the trajectory looks like archetypical *overcoming*. But what is balancing the evident force of gravity? What is it that gravity is overcoming on the downward path? Consider also the upward trajectory which, interpreted as a nonviolent *continuous motion* (result), needs a continuous force cause. In other words, the problem context cues a number of schemes that all have one missing element, an element that some sort of upward force residing within the ball could occupy. That force, if it *died away* (another natural phenomenon) would solve the problems posed. It would be the conflicting partner of gravity that, while greater than gravity, would propel the object upward, while equal to gravity, would balance it at the peak, and as it decayed, would leave gravity to overcome it.[1]

What I am saying is that something like impetus is an invention particular to this or some relatively small class of problems rather than a fundamental theoretical construct of intuitive physics. To be sure, the idea shows a significant resonance with many elements of intuitive physics; it can essentially be derived in the context of a toss from the set of fragments enumerated above. Yet, as I will indicate below, it hardly has the priority that would mark it as the core of a systematic, theoretical view. Instead, I claim that understanding intuitive physics necessarily means understanding the kind of pieces into which I have just decomposed impetuslike explanations.

Showing a decomposition of impetuslike explanations into a set of plausible pieces that do not include the notion of impetus is one piece of evidence undermining the impetus theory theory. But we must look to a broad range of circumstances to really prove the case. We should find the following phenomena that distinguish the impetus theory from the p-prims theory:

1. Because the impetus theory is a pattern that emerges from more invariant pieces, we should see those pieces in other contexts, either alone or in other combinations with no hint of impetus.
2. Indeed, the list I have presented is hardly exhaustive, according to the p-prims view, and it should not be hard to find situations in which subjects give reliable responses that have nothing whatsoever to do with impetus or any of the p-prims in the short list I presented here.
3. We should expect a significant spread in answers subjects give to problems like the toss. Even if they all have the same set of p-prims at the base of their intuitive knowledge, we should not expect that they all uniformly derive an impetuslike explanation in this context. Many of these should involve the same or similar p-prims, but in other combinations and attached to other features of the problem.
4. We need to look in some detail at the dynamic of the generation of impetuslike explanations, even if they reliably occur. There should be an

[1] It is worth noting an additional consonance between impetus and naive p-prims. The force residing in the ball is a particular manifestation of the kind of animism that Piaget describes in some children's descriptions of physical events. A moving ball exhibits an independent motion and can even exhibit agency in making other things move in collisions. Yet it is not alive; it cannot initiate its own motion. Because it does not originate in the ball, children see the ball's agency in terms of something that is *transferred* to it. Naive p-prims having to do with substance and transfer approximate the state of affairs in description by reifying the quality of motion transferred (roughly, its direction and magnitude), as a restricted kind of life. In children, we might call this "animism." Physicists call it "momentum."

observable genesis; impetus should not emerge instantly and fully. There should be "waffling" and uncertainty and no great commitment to the impetus idea. The notion should be somewhat unstable even if it is generated, and should occasionally give way to other kinds of descriptions and explanations that may be consistent with the above or an extended list of p-prims, but not with impetus as a guiding notion.

Even though I have not personally studied the toss problem in detail, I do not think it is very hard to find evidence for 3 and 4 above. In the best of cases, one finds explanations smacking of impetus occurring in less than half the subjects, even by McCloskey's own reckoning. Because it is almost impossible to find subjects to admit a strong commitment to any particular interpretation in a problem context, let alone to that interpretation as a general fact that determines a fundamental law of nature, waffling, alternate explanations, and so forth, are really the order of the day in essentially any intuitive physics protocol.

Instead of pursuing those lines, I will provide some examples in categories 1 and 2 above, of other contexts where one sees impetus-fragment p-prims without impetus, and where one finds situations governed by p-prims having nothing to do with impetus or even the list of p-prims given above.

It turns out that it is terribly easy to find situations that do not generate impetus explanations. In fact, if one asks for a description or explanation of what happens when one simply drops a ball from rest, the second half of the toss, one almost never sees impetus responses. From the p-prims view, this is easy to explain. There is no conflict, there is no balancing, there is nothing, in fact, that needs another agency like impetus to explain it. Gravity simply gets what it wants, the ball falls down without the aid of any internal impetus.[1]

In this regard, it is useful to look at the history of science, at one who was really trying to build theories. Galileo took a look at impetus as an explanation of the toss at one stage in his dialogues. He saw clearly the need to find a way to extend that analysis beyond the initial context that suggested it in order to test and develop the notion into a genuine theory. After looking at the toss, he developed a clever, but hardly intuitively compelling, explanation of the fall of a dropped object using impetus. But the intuitive thinker does not develop sophisticated arguments to extend the scope of a notion, neither does he even feel the need to talk about tossing and dropping in uniform terms.

Galileo provides a good study of intuitive versus theoretical knowledge systems in other respects. He systematically highlighted intuitive arguments that

[1] McCloskey explains this particular restriction in domain of application of the impetus theory by asserting it: impetus does not apply to situations of carrying. Besides being subject to the question, "Why?" my conjecture is that if one studied many problems at the edge of the common range of use of impetus ideas, one would need many such unmotivated patches. Theory theories, in general, will not provide a small enough grain size to cope with people's commonsense physical ideas. The problems that follow in the text continue to make this point.

counter his ideas as part of his expositional technique and, one by one, defused them. One can not only build a compelling list of p-prims from this (to appear), but again one sees how hard it is to build any core of ideas into a theoretical coherence that extends beyond a very limited context. At every turn, intuition suggests other ways to think about the situation. Even more impressive is how many contexts and arguments it takes to clarify and make precise even the core ideas. Galileo runs through at least one-half dozen intuitive frames of analysis simply to explain and make plausible the single idea that an object dropped from rest accelerates uniformly, which idea he took to be one of his greatest accomplishments. For the historically interested, I highly recommend rereading the relevant Galileo in the light of these issues (Galileo, 1954, p. 165, and surrounding).

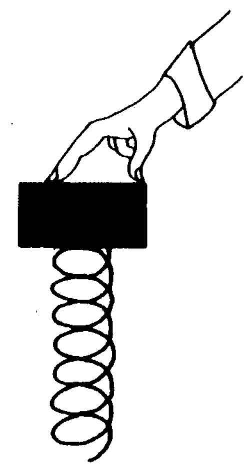

FIGURE 4.1. How far must one compress a spring in order for it to toss an object into the air?

Figure 4.1 shows a problem that cues certain p-prim fragments, but no impetus. The question is: How far must one press a brick down onto a spring, compressing it, before the spring will be able to toss the brick into the air when one releases the compression? Students frequently see this as a straightforward situation of dynamic balance and overcoming. If you compress the spring until it is pressing upward harder than gravity is pressing downward, the spring will overcome gravity on releasing your hand, throwing the brick into the air. If one writes down the equation expressing this condition, it turns out one has only specified the equilibrium point where the spring is compressed enough to support the brick unaided. Instead, the problem *really* is a question of impetus from the physicist's point of view; can the spring provide a net positive impetus (momentum) to the brick by the time it has completely extended itself? It is not enough that the spring supplies a force greater than gravity at the beginning of the "toss."

(a)

FIGURE 4.2. A boxcar in space whizzing past a planet. What path does the mass on the spring take?

Figure 4.2 shows an equally striking situation that does not provoke impetus responses even though, in this case again, impetus responses would be more appropriate than the p-prims that are actually cued. A mass is attached to a spring, which, in turn, is attached to the base of a boxcar (and constrained so that it cannot flop over). The boxcar happens to be running at a constant and huge speed along a perfectly straight railroad track in space. As the boxcar passes a planet (and the mass feels the gravitational pull of the planet) what path does the mass travel along?

(b)

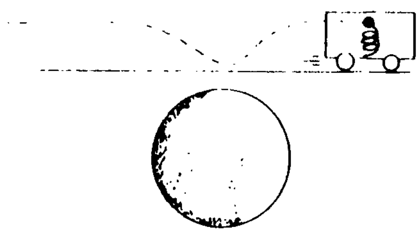

Figure 4.2b is an enormously attractive answer. Not only is the answer visually appealing, but everyone knows that springs compress an amount that is somehow proportional to the force on them. As the force gets symmetrically stronger and weaker while the boxcar flies by the planet, the spring's compression should follow suit.

(c)

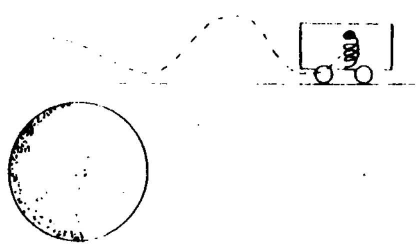

Figure 4.2c is the correct answer. The mass acquires a downward impetus (momentum) from the pull of the planet, which carries it downward even after the closest approach to the planet. Thereafter, the mass will oscillate on the end of the spring as the gravitational perturbation dies out.

Students show an amazing resistance to thinking about this problem dynamically, in terms of impetus, momentum, or any other dynamic concept. Instead, it seems so obviously a question that can be answered by what I call the "spring scale" p-prim—squishy things compress an amount proportional to the force on them—that often not even an hour or more of probing that highlights the dynamic nature of the problem can shake the conviction of even bright MIT students that "spring scale" is the right way to think about it.

Circular Motion

I would like to report briefly on another context in which p-prims can be compared to the theory theory, and to the impetus theory in particular. McCloskey claims that rotational problems are handled by a branch of the impetus theory that holds that there also exists a second kind of impetus, circular impetus, that works in basically the same way as linear impetus. The difference is only that circular motion is involved rather than straight-line motion. Circular impetus acts to continue circular motion after a circular push has stopped.

Rather than put this idea to a test at the extremes of its predictive power, Tamar Globerson (from the School of Education, Tel Aviv University) and I decided to test it with only very small variations of a single problem. Again, the point is not whether subjects ever give any indication of circular impetus explanations, but whether impetus is any sort of complete characterization of or even a reasonable approximation to people's knowledge state as far as circular motion is concerned. We used variations on Piaget's sling problem: What happens to a ball twirling around on the end of a string if one cuts the string? The important thing is that, from the point of view of Newtonian mechanics, and even from the point of view of the impetus theory, the problems are all identical: What happens when a circular motion is aborted by removing the circumstances causing it? A physicist or one who holds an impetus theory should see through the small variations and give a uniform response.

In addition to the ball-and-string form of the problem, we used the following variations, among others:

1. A ball is moving in a circular tube. What happens when it leaves the tube? Suppose one adds a short circular extension to the tube, then removes it just as the ball is about to enter it? Does anything change?
2. A ball is moving on a table under a circular tunnel. If the tunnel is removed, what happens?
3. Instead of asking subjects to predict the results, what happens if subjects are asked merely to react to the plausibility of various paths presented in a computer simulation of the cut-string problem?

4. Suppose one asks subjects to rethink their responses on the basis of suggestions that "some other people think . . ." or are merely prompted to remember particular phenomena related to circular motion?

In a preliminary study that involved early elementary school children (around first grade), late elementary school children (around fifth grade) and "physics naive" adults, we found striking variability in the answers and justifications that subjects gave. Subjects frequently gave multiple kinds of predictions and explanations, and these answers changed according to the circumstances of presentation. Literally no one gave and maintained a pure circular impetus explanation. A particularly strong example of the variability came with respect to the issue of centrifugal force. Whereas "a pull to the outside" (not necessarily known by the name *centrifugal force*) was perceived in the ball-and-string presentation, few thought the same thing happened inside a tube. The explanation for this seems to have something to do with the overt and focused "tug to the center" offered by the string as opposed to mere "guidance" offered by the tube. Few of our subjects were sophisticated enough to spontaneously see a force acting inside the tube.

Mild interventions aimed at prompting remembrance of centrifugal phenomena—"Do you remember what you feel when you drive rapidly around a corner in a car?"—frequently caused subjects to shift predictions. In particular, the prediction that the ball should move outward at a 45° angle to a tangent to the circle became quite salient after such an intervention. The explanation subjects gave for this prediction is that the ball feels two tendencies: one to keep moving forward (either straight or in a circle) and one due to the outward pull. Thus it will actually follow a compromise, the 45° path. Note how far such prediction is from a simple, exclusive adherence to a circular impetus that has no provision for centrifugal force or for any principles of combining multiple influences in circular situations. Although not a single subject in our study gave this answer spontaneously, more than a quarter of the subjects who were prompted to think about centrifugal force declared they preferred this prediction to all other predictions, either self-produced or offered by us, at the end of the session. In view of such data, it is evidently an extreme oversimplification, at least, to attribute theorylike status to circular impetus responses.

One of the most striking effects noticed in this experiment, and one which needs follow-up study, has to do with modality of presentation. First, it must be noted that the youngest and oldest subjects gave quite different predictions, explanations, and justifications. For example, a significant percentage of the younger subjects said the ball would stop after it left the tube or the string was cut "because it wouldn't know which way to go," whereas none of the older

elementary school students or adults gave such a response. In contrast, when asked merely to rate the plausibility of various results depicted in a computer simulation, essentially all age differences disappeared. It is as if a visual and inarticulate ability to judge plausibility of various motions develops quite early whereas more articulate explanations and theoretical constructs continue to evolve substantially.

PUTTING THE PIECES TOGETHER

I have motivated a view of intuitive knowledge in physics that poses fundamental educational problems in terms quite different from disabusing students of a theory competitive to Newton's. Indeed, perhaps the most fundamental problem is the simple fact that students come to physics classes with no theory at all, but instead are used to dealing with the world on a catch-as-catch-can basis, where it is quite fair to change tactics whenever the problem is minutely varied. There is an entirely different style of thinking involved when one comes to the stage that, for example, all the circular motion problems that Globerson and I cooked up are perceived and acted on as trivial variations of the same problem. I firmly believe that students who can articulately espouse *any* systematic view of the physical world would be far better prepared for physics courses than those who can be coaxed into reciting the right words, yet behave as if every new problem were an occasion to invent another explanation. This section is aimed at commenting on the uses of computers in education that are consonant with "knowledge in pieces."

Microworlds

Let me begin by going back to an old idea. When we ask ourselves where people get the funny ideas that they have, we must, like Piaget, look to experience—not just experience as judged from an abstract view of what people are doing, but experience as felt internally, as judged by the extent to which people discern structure in the experience, and to the extent that this structure is contributing to the development of new mental structures. The key here is continuity. We cannot expect to have students learn things radically distant from their current state of understanding. Nor will they learn things that have a radically different character, such as the extremely systematic view of the mechanical universe provided by Newton, except by progressing through stages of understanding that, by degrees, approximate the final state.

So we should begin with experiences that have roughly the same character as those that generate and support intuitive physics as we find it.[3] This is the idea of microworlds, constructing artificial realities that intersect enough with students' ideas that they can immediately begin to manipulate them, but whose "deep structure," if you like, leads inevitably beyond those initial perceptions and conceptions. I and others have talked and written much about this idea in other contexts (see, for example, diSessa, 1982, and Papert, 1980), so I won't belabor it here, except to note that computers are so flexible as a design material that we should soon see, if we have not already begun to see, a boom in materials-based, experiential learning. Computers are so versatile in crafting interactive environments that we are more limited by our theoretical notions of learning and our imaginations. We can go far beyond the constraints of conventional materials, which are limited to an interaction of "push, pull, poke, and position" in a high-friction universe.

There are two particular ideas extending the notion of microworlds that are appropriate to the discussion of fragmented knowledge. These are relatively new ideas for me, or better, ideas whose importance has only gradually impressed me as I have come to see more of the fragmentation of intuitive knowledge and the educational problems that it poses. Rather than being closed-form ideas, these are more heuristics to help guide the design of microworlds.

The first is what I will call *mega-microworlds*. Simply put, the idea is that a single perspective is almost never enough to build a well-integrated and widely applicable understanding of the sort that we would want to call "scientific." The most carefully crafted experience just won't do it by itself. Instead, one needs to build clusters of these so that students can become involved in many ways over an extended period of time. Even from my own practical experience in building microworlds, it seems we must have quite limited expectations for any particular one, but must turn to building a multiplicity of them with a common thread. A bit metaphorically, we must find proper contexts to express all of the right collection of perspectives, the right set of p-prims that can be integrated into a new scientific understanding. This idea is developed in diSessa (1986).

The second extension of the microworld idea I call *textured microworlds*. This is an optimistic position that we are almost in a position scientifically to note, perhaps not one by one, but at least by the general class, the p-prims that we expect students to be abstracting and recombining in their microworld experiences. This amounts to saying we can nearly chart the essentials of experiential knowledge. Instead of just designing an activity or set of activities, hoping for the best, we should begin to have more precise expectations about what exactly students will learn from various contexts. Such microstructure should allow us to do more than see success or failure when we have designed and built a microworld and let a student go off for a while to play in it. We should begin to expect to be able to assess partial successes, to "debug," if you like: find the places where knowledge must be patched. In microworlds that come with such a rich, theoretical texture, we should be much more capable of making principled interventions. The beginnings of such a texture for a microworld for learning Newtonian mechanics are presented in diSessa (1982).

The concepts of mega- and textured-microworlds, of course, need not be limited to computer-based materials. In fact, Marlene Kliman (presently a graduate student at Harvard) and I have tried to chart in a much more refined way than has until now been attempted the intuitive knowledge that becomes involved when children interact with a relatively common and nonelectronic piece of pedagogical instrumentation, the balance scale. The result is a data base that we hope will add tremendously to a teacher's ability to watch a child and know what is going on, to know what knowledge is being used, where it comes from, and where it might go developmentally (Kliman, 1987). The data base is an excellent place to accumulate profitable interventions.[4]

[3]Here, by the way, we find another shortcoming of the theory theory. If theories are the stock and trade of naive understanding of the physical world, surely it is important to ask how these theories develop. What replaces in people the historical and social forces that create theories in public science? In contrast, at least some parts of the beginnings of the phenomenological physics that I have described are comparatively unproblematic. People make many relatively simple observations about the physical world, keep them, for the most part, relatively isolated from one another, and only gradually percolate some of these up to a level at which they can even be surprised that these phenomena do not hold in some circumstance or another.

Mediating Between Formalisms and Experience

Computers have a very special niche at the interface between, on the one side, formalisms—those grand unifiers of science where one can write down "F = ma" and summarize all of Newtonian Mechanics in a little box—and, on the other side, experience with its apparently infinite fragmentation. A major problem with formalisms in past pedagogy is that they have stood quite apart from intuitive knowledge. Indeed, they are often made to be the antithesis to intuitive ideas, rather than to be productively engaging of them. The computer can play a multitude of important roles squarely between these poles, making for productive transitions in both directions. Generally, programming and computer modeling can profitably interject formalisms into an otherwise experiential microworld. But I would like to give more specific examples that I think are indicative of the richness and importance of this way of viewing the role of computers in science education.

[4]Even in such a context, computers may play a significant role in helping teachers keep track of their students' knowledge state and in helping to suggest interventions out of the large data base of intuitive knowledge. Advances in artificial intelligence should eventually allow us to entirely automate this feedback loop

The first example is a microworld I designed a number of years ago for optics. The core of the microworld is a "simulation" that allows one to place a number of optical elements of different kinds, lenses, mirrors, prisms, and so on, into a field, and then shoot rays singly or in clusters through the constructed optical system. This is an automation of a little formalism called *ray tracing* that was developed in order to help people think about optics. From the way the rays travel through the system, one can figure out all other optical properties of the system.

Unfortunately, I discovered that, for the most part, this microworld was a dud. Students would sit down and play with it for a while, and they would happily do exercises I assigned to them on it—it is a much better tool than paper and pencil—but an important ingredient was missing. The system simply does not have any of the experiential feel of optical phenomena in the real world. The most immediate consequence of this was that students were not motivated to play and try things out, but instead treated it like the formalism it really was, as a tool to be trotted out when absolutely forced to.

Now, I think I know how to fix this microworld. A student of mine, Ed Lay, added a single feature to the system which has entirely changed the feeling of the system. The feature is that in addition to placing optical objects down, one can place *things to see* and can ask the system to show you what is seen from any vantage point. All of a sudden abstract questions become experiential and immediate. We had the experience within our own research group. (I have not yet had the chance to actually try this modified microworld with students.) You just poke around a bit and all of a sudden, "Grief, why is the picture upside down now?" Or, "Why is it in or out of focus?" Design criteria, magnification, lack of distortion, become directly observable. And the formalism, ray tracing, is always available in the same context to supply precisely what it is best at, careful analysis. The formalism is seen as a powerful tool, not to be mistaken for the object of study itself. Building analytic or other formal tools right into experiential environments should become more and more a standard part of microworld design.

Intuitive understanding of traditional formalisms is hard won because the phenomenology of operating with them is so subtle, so far away from daily experience and naive p-prims. Could it be that with computers we can design more dynamic and interactive formalisms that transcent this problem? The idea is not to juxtapose experiential and formal points of view as above, but to fuse them. One seeks to build things that are understandable and engaging in their own right, so that students have a painless and productive engagement with them, but things that happen also to be good formalisms for representing and thinking about a broad range of phenomena. I call manipulable systems that can serve as general and precise formalisms but which retain for students a sense of familiarity and evident controllability *semi-formalisms*. Semi-formalisms will

often take the form of construction kits to make things that can be viewed either as toys or as formal models.

Both of the following are examples of flow systems in which some "abstract" stuff flows from node to node according to rules that specify how it should flow. In fact, both of these might be constructions in an as-yet only partially implemented construction kit that allows many types of flow systems to be simply generated. Stuff moving around from place to place is understandable and interesting in its own right, yet happens also to be a powerful and extendible way of modeling many real-world situations that may or may not literally involve flow. I have for years advocated flow as the core to a rigorous but intuitively accessible way of thinking about many areas. (See, for example, diSessa, (1980). Computers can build on this initial accessibility by providing more precise means of control and analysis in what remains familiar and experiential.

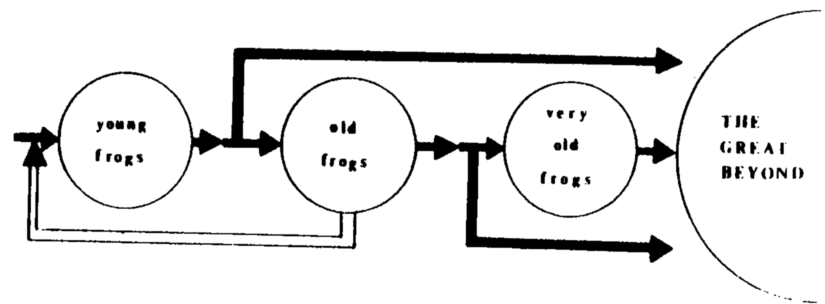

FIGURE 4.3. A flow system used to model a simple ecological system.

Figure 4.3 shows a mock-up of a flow system used as a model of a simple ecological system. Each of the nodes represent the number of frogs at various age levels. Numbers or some thermometer-type analog display can show how many frogs occupy each node. On the left, one has young frogs. In the middle are older frogs, and on the right are very old frogs. Now, the connections show the routes that frogs can take as they age. Some young frogs grow into older frogs. But some, because of famine or disease, go to "the great beyond." Older frogs have the same options. The oldest frogs have no choice, but go directly to the great beyond. There is a single control line in the system as shown; the number of young frogs entering depends on the number of older frogs and is independent of the number of young and very old frogs.

Even with such a simple setup, one can see interesting phenomena. For example, if one artificially sets up a surge in births at one stage, a wave in populations will ripple through the system, and secondary waves will be generated as each baby boom reaches the fertile age.

Things get even more interesting if we introduce, for example, the number of flies in the ecological system. It should control the branching ratio of frogs that go on to older age levels compared to frogs that go to the great beyond. If we add a negative influence between the total population of frogs and the number of flies (more frogs eat more flies), then one can get extremely intricate patterns of behavior.

What makes this an interesting environment is that it has independent layers of understandability: First, as a system that moves "stuff" around according to a simple set of influences and controls, one can set up and play (perhaps literally as games) with systems with few parts, all of whose interactions are simple, but, in the large, exhibit complex and sometimes subtle behavior. Just as well, one can use this as a little formalism to experiment with various realistic and unrealistic models of actual physical systems, of which the ecology system is only one example.

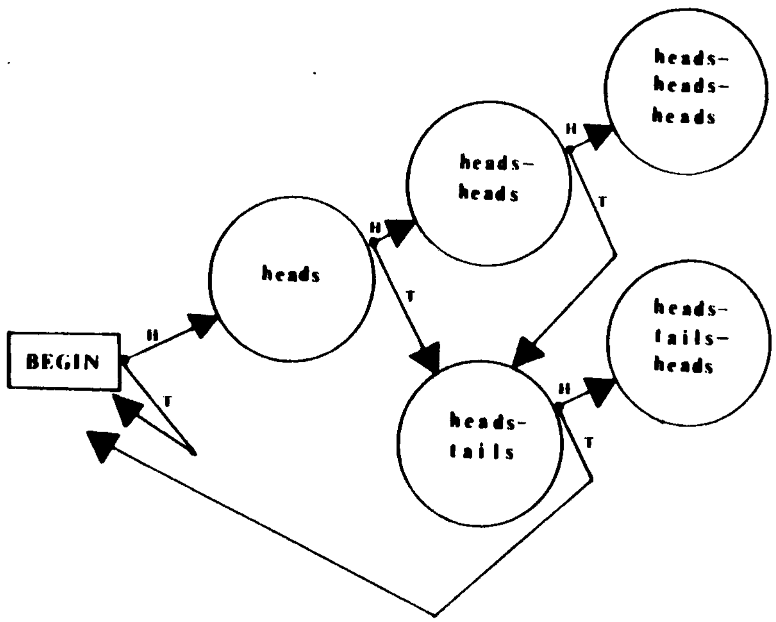

FIGURE 4.4. A marker "flows" along a state transition graph, showing progress toward two goals.

Let me show another example of a very different modeling function built out of the same formalism. Figure 4.4 shows a system that solves a problem in probability theory relating to coin tosses. Consider a question such as: What can you say about the likelihood that you will toss a sequence of heads-heads-heads before heads-tails-heads? Figure 4.4 depicts what is called a "state transition graph" and works as follows: Suppose you first toss heads. Then you can move a marker, which shows the state of your progress toward the two goals, from

"begin" to the "heads" node. If you subsequently toss heads or tails, you would move your marker along the appropriately marked path to "heads-heads" or "heads-tails" respectively, moving closer to one goal or the other. If you toss in sequence heads, tails, and then tails, then you must follow the path marked T from "heads-tails" back to "begin," because you are nowhere near either of your goals and must effectively start from scratch. On the other hand, a tails after heads-heads puts you in position to reach the heads-tails-heads goal with a subsequent heads toss, which is equivalent to having tossed heads-tails. (In this case, the initial heads is no longer relevant.)

Of course, one could make this little system with paper and pencil, and move markers around, but imagine how much easier it is with a computer system designed for such things. Imagine, as well, children have already encountered the flow system, perhaps with frogs, so that they know how to build their own systems, set them up, make them run automatically, and so forth. They can start a million coins at "begin," then watch them all move from node to node, one-half taking each path from each node on each toss. Or one could take a single coin and run it a million times through the system, taking the heads or tails path out of a node each time by chance. Incidentally, although it is no more likely that you will throw a heads-tails-heads than a heads-heads-heads in three coin tosses, you have a much better chance of reaching a heads-tails-heads goal before a heads-heads-heads. This should be apprehensible from the flow pattern represented in Figure 4.4, even without running it.

All in all, the flow kit should serve as an interactive formalism for modeling many sorts of systems. What makes it a semiformalism is that the "formalism" is itself something that is comprehensible and something with which one can immediately play because it is based on a simple metaphor, movement of things.

Intuitive Epistemology

I would like to mention a final area of leverage in using computers to attack the fragmentation one finds in intuitive knowledge systems. I have done some preliminary study of an area that I call "intuitive epistemology." This is another intuitive knowledge system that, however, concerns the phenomenology of personal intellectual functioning rather than phenomenology of the physical world. People have perceptions about what happens, about what causes what, about what is important and what is not concerning knowledge, its development, and its deployment. In some cases these ideas also seem to be almost theoretical, but the same caveats are warranted here as with intuitive physics.

In diSessa (1985) I develop two case studies of students with remarkably contrasting intuitive epistemologies. One of these looks in many instances to be similar to my own p-prims epistemology, particularly with respect to the relation of intuitive knowledge and textbook physics. The other might be caricatured as a theory that physics resides only in the equations and formalisms and that intuitive knowledge is only so much confusion. One would guess that the second student might be at a distinct disadvantage in learning physics with such a poor model of what it is he is trying to learn. Indeed, this appeared to be the case to the limits of this small study.

One would really like to walk up to the second student (and, in my experience as a physics instructor, there are many like him) and say, "Look, you have a fragmented, piecemeal jumble of an intuitive knowledge system, and that's fine; it contains a lot of the right pieces. You must find, cultivate, and refine these pieces. And, above all, you really need to concentrate on integrating that system. Science is, after all, at least as well characterized by its systematicity as by its content." Of course, that is bound to fail if you do it like that. That is the confrontational approach of the theory theory. Even if such students could understand literally what was said, even if they believed you, they still need to have some sense for what it feels like to use intuition properly or to be scientifically coherent; they need to have some sense for what one does to unify.

Again, because confrontation will not work does not mean the battle is lost. Again, computers can be extraordinarily helpful in providing students with experiences that meet their intuitive ideas and develop them toward more integrated and profitable points of view. Here, I will make a short list of techniques. Details appear in diSessa (1985).

Amplifying Qualitative Phenomena. Much science education is dominated by analytic methods. Problems are posed in terms of known and unknown quantities. So problem solving is often perceived by students as finding the right equation, which represents knowledge to them, and grinding through a little arithmetic or algebra. Instead, their images of knowing and learning should be different if they deal constantly with phenomena at the qualitative level. This is microworlds at the meta-level. Experimentation and research on real-world systems could serve this role, except we have much greater difficulty designing such experiences to have the important ideas represented near the visible structure and in the direct principles of manipulation of the system. Semiformalisms and related computational devices have a role here too in providing nonanalytic but still technical help in solving problems.

Undermining Naive Realism. If analytic methods provide poor models of knowledge, the absence of salient alternatives can lead to assumptions that knowing (after learning) is simple and direct, even if problem solving is not. What one would like to have are explicit schemes that represent the world in a powerful way, yet which have a visible human genesis and must be used carefully. Indeed, the notion of representation—of a schema that one builds, but which accurately reflects some reality without being mistaken for it—is central and visibly pervasive in computational systems. Every time one builds a simulation, one selects or invents a representation of reality, discovers it is not quite right, throws it out, and rebuilds or invents "patches" to overcome inadequacies. This can be a much better process from which to abstract ideas about knowing than always having been supplied with a thoroughly debugged and apparently simple law.

Changing the Quality and Time Scale of Exercises. One of the great experiences of my scientific career was to discover in high school that I could profitably think about a problem for weeks, that I could get valid insights, make progress, see things gradually fall into place, rather than just "find the solution" like a needle in a haystack. Yet I had this experience in none of my classes, but on an exam for a summer science program for which I was applying. Indeed, the vast majority of work done by students in school classes is of the 20-minutes-or-less-per-problem type. This is hardly fertile ground to promote awareness of learning processes that may be months or years long. Instead, I believe we must engage students in research and design in many ways more typical of professional practice than schoolish exercises. In the same way as computers have become indispensable tools for scientists and engineers, they are nearly essential here. Engaging in research and design is not easy, and without significant help in terms of good areas to research (microworlds), good tools with which to do analysis, and good material to design with, it may be nearly impossible for students. In my years of teaching research and design courses to high school students and undergraduates, not a single student has *not* used the computer, even though I never hinted that it was necessary, and I even occasionally encouraged students to work with materials other than the computer. This is due to more than the fact that computers are fun. For certain things, their utility is more than obvious to everyone.

Changing Subjects That Are Taught. Intuitive epistemology changes our pedagogical agenda. We are not only concerned with the "stuff" that students learn, but the process that they go through, and the meta-cognitive abstractions that they make from that experience. Some things to learn may be significantly more or less attractive from this point of view. Indeed, some of the best things to learn from this point of view may be unteachable without computers. Some readers may know the book I wrote with Hal Abelson on Turtle Geometry (Abelson & diSessa, 1981). Anyone who does cannot doubt that teaching that material would be difficult or impossible without the computational experiences built into the book. Learning computational geometry without touching a computer would be like learning physics never having touched, pushed, or pulled a physical object. I cannot begin to try to convince that learning Turtle Geometry is good for your intuitive epistemology here, so I refer the interested to the book and to things that have been written about it (e.g., diSessa, 1979). I am not sanguine about when the educational establishment can accept such radical changes at the core of its sacred curriculum, but computers can at least be an edge in.

A CLOSING NOTE

The fragmented system of intuitive knowledge that we find in our students poses significant educational problems. We need not throw up our hands, however, but should roll up our sleeves and get to work with the best tools we have. Computers are such tools.

With all my optimism about computers, I must emphasize that they are not magical instruments to engage and integrate intuitive knowledge. They will not help independent of what we do with them. We need not design microworlds or attempt to devise semiformalisms. We certainly need not change the subject matter we teach with them. Indeed, the first guesses for how we should integrate computers into education did none of the profitable things I listed here. One can easily teach the same old things with the same old 20-minute exercises. We could easily undermine qualitative methods, and emphasize numbers and other formal methods so that the computer is really just a number or symbol cruncher. The choice is ours whether computers will help solve or aggravate the problem of knowledge in pieces.

REFERENCES

Abelson, H., & diSessa, A. (1981). *Turtle Geometry: The computer as a medium for exploring mathematics*. Cambridge, MA: MIT Press.

diSessa, A. (1979). Learnable representations of knowledge. In J. Lochhead & J. Clement (Eds.), *Cognitive process instruction*. Philadelphia: Franklin Institute Press.

diSessa, A. (1980). Momentum flow as an alternative perspective in elementary mechanics. *American Journal of Physics, 48* (5), 365-369.

diSessa, A. (1982). Unlearning Aristotelian physics: A study of knowledge-based learning. *Cognitive Science, 6* (1), 37-75.

diSessa, A. (1983). Phenomenology and the evolution of intuition. In D. Gentner & A. Stevens (Eds.), *Mental models* (pp. 15-33). Hillsdale, NJ: Lawrence Erlbaum Associates.

diSessa, A. (1985). Learning about knowing. In E. Klein (Ed.), *Children and computers: New directions for child development (28)*, (pp. 97-124). San Francisco: Jossey-Bass.

diSessa, A. (1986). Artificial worlds and real experience. *Instructional science, 14* (3 and 4), 207-227.

Galileo Galilei (1954). *Dialogs concerning two new sciences* (H. Crew & A. deSilvio Trans.). New York: Dover.

Kliman, M. (1987). Children's learning about the balance scale. *Instructional Science, 15* (4), 307-340.

McCloskey, M. (1983, July). Intuitive physics. *Scientific American*, pp. 122-130.

McCloskey, M. (1983). Naive theories of motion. In D. Gentner & A. Stevens (Eds.), *Mental models* (pp. 299-324). Hillsdale, NJ: Lawrence Erlbaum Associates.

Papert, S. A. (1980). *Mindstorms: Children, computers and powerful ideas*. New York: Basic Books.

Chapter 5

Some Pieces of the Puzzle

Jack Lochhead
University of Massachusetts

Ten years ago Jean Piaget spoke to the Piaget Society on the topic of "Correspondences and Transformations." As one ought to expect from any great mind addressing a society named after himself, Piaget began by questioning the basic assumptions of his own theory. He indicated that although he had originally believed that mental processes could be understood solely in terms of transformations, he no longer believed this to be the case.

Piaget stated that the use of correspondences is a necessary precursor to the ability to think in terms of transformations. One reason for this is that transformations require mastery of some form of reversibility, whereas correspondences do not. Two objects can be exactly alike, indistinguishable, but they cannot be absolutely different, thus there is no inverse correspondence. But even though correspondences appear to be simpler than transformations, the issue of how they can be performed is far less easy to ascertain. To determine difference, one need only find a comparison that fails; to determine sameness, one needs to perform possibly an infinite number of comparisons.

Although it may well be impossible to determine true identity, it is fairly easy to create vague similarities based on sloppy criteria. This roughly, is what Vygotsky's (1962) preconceptual thinker does when he formulates heaps. It is also similar to what Rosnick (1982) referred to as the "generation of undifferentiated conglomerates."

I believe there are some interesting correspondences between these ideas expressed by Piaget, Vygotsky, and Rosnick, and those in diSessa's chapter (chapter 4) on knowledge in pieces. In all cases there is the suggestion that before knowledge can be organized in comprehensive global structures it first must be collected piecemeal. Coherence and self-consistency are not possible in