CHAPTER 2

Cognitive Principles and Guidelines for Instruction[1]

He who loves practice without theory
is like the sailor who boards ship
without a rudder and compass
and never knows where he may cast.
Leonardo da Vinci
quoted in [Fripp 2000]

When we present instruction to our students, we always build in our assumptions: our expectations as to what students will do with whatever we give them, our assumptions about the nature of learning, and our assumptions about the goals of our particular instruction. Sometimes those assumptions are explicit, but more frequently they are unstated and rarely discussed. Some pertain to choices we get to make, such as the goals of instruction. Others are assumptions about the nature and response of a physical system—the student—and these are places where we can be right or wrong about how the system works.

If we design our instruction on the basis of incorrect assumptions about our students, we can get results that differ dramatically from what we expect. To design effective instruction—indeed to help us understand what effective instruction means—we need to understand a bit about how the student mind functions. Much has been learned about how the mind works from studies in cognitive science, neuroscience, and education over the past 50 years. In this chapter and the next, I summarize the critical points of the cognitive model and organize the information in a way that relates to the instructional context. I then consider some specific implications for physics instruction: the impact of considering students' prior knowledge and the relevant components of physics learning other than content. The chapter ends with a discussion of how our explicit cognitive model of student learning can provide guidelines to help us both understand what is happening in our classroom and improve our instruction.

---

[1]This chapter is based in part on the paper [Redish 1994].

THE COGNITIVE MODEL

To understand learning, we must understand memory—how information is stored in the brain. Modern cognitive science now has complex and detailed structural information about how memory works. For some simple organisms like the marine snail *Aplysia*,[2] the process is understood down to the level of neuron chemistry [Squire 1999]. We don't need that level of detail for the "application" of understanding physics teaching and learning. A few simple principles will help us understand the critical issues.

Models of memory

It is clear, from all the different things that people can do that require memory, that memory is a highly complex and structured phenomenon. Fortunately, we only need to understand a small part of the structure to get started in learning more about how to teach physics effectively. There are a few critical ideas that are relevant for us. First, memory can be divided into two primary components: *working memory* and *long-term memory*.

- Working memory is fast but limited. It can only handle a small number of data blocks, and the content tends to fade after a few seconds.
- Long-term memory can hold a huge amount of information—facts, data, and rules for how to use and process them—and the information can be maintained for long periods (for years or even decades).
- Most information in long-term memory is not immediately accessible. Using information from long-term memory requires that it be activated (brought into working memory).
- Activation of information in long-term memory is productive (created on the spot from small, stable parts) and associative (activating an element leads to activation of other elements).

In the rest of this section, I elaborate on and discuss some of the characteristics of memory that are particularly relevant for us.

1. Working memory

Working memory appears to be the part of our memory that we use for problem solving, processing information, and maintaining information in our consciousness. Cognitive and neuroscientists have studied working memory rather extensively. Not only is it very important to understand working memory in order to understand thinking, but working memory can be studied with direct, carefully controlled experiments [Baddeley 1998]. For our concerns here, two characteristics are particularly important:

- Working memory is limited.
- Working memory contains distinct verbal and visual parts.

[2] *Aplysia* has a nervous system with a very small number of neurons—about 20,000, some of them very large—and a very simple behavioral repertoire. As a consequence, it is a favorite subject for reductionist neuroscientists. See, for example, [Squire 1999].

Working memory is limited. The first critical point about working memory for us to consider is that working memory can only handle a fairly small number of "units" or "chunks" at one time. Early experiments [Miller 1956] suggested that the number was "7 ± 2". We cannot understand that number until we ask, "What do they mean by a unit?" Miller's experiments involved strings of numbers, letters, or words. But clearly people can construct very large arguments! If I had to write out everything that is contained in the proof of a theorem in string theory, it would take hundreds, perhaps thousands, of pages. The key, of course, is that we don't (write out everything, that is). Our knowledge is combined into hierarchies of blocks (or chunks) that we can work with even with our limited short-term processing ability.

You can see the structure of working memory in your own head by trying to memorize the following string of numbers:

3 5 2 9 7 4 3 1 0 4 8 5

Look at it, read it aloud to yourself, or have someone read it aloud to you; look away for 10 seconds and try to write the string down without looking at it. How did you do? Most people given this task will get some right at the beginning, some right at the end, and do very badly in the middle. Now try the same task with the following string

1 7 7 6 1 8 6 5 1 9 4 1

If you are an American *and* if you noticed the pattern (try grouping the numbers in blocks of four), you are likely to have no trouble getting them all correct—even a week later.

The groups of four numbers in the second string are "chunks"—each string of four numbers is associated with a year, and is not seen as four independent numbers. The interesting thing to note here is that some people look at the second string of numbers and do not automatically notice that it conveniently groups into dates. These people have just as much trouble with the second string as with the first—until the chunking is pointed out to them. This points out a number of interesting issues about working memory.[3]

- Working memory has a limited size, but it can work with chunks that can have considerable structure.
- Working memory does not function independently of long-term memory. The interpretation and understanding of items in working memory depend on their presence and associations in long-term memory.
- The effective number of chunks a piece of information takes up in working memory depends on the individual's knowledge and mental state (i.e., whether the knowledge has been activated).

This second point is fairly obvious when we think about reading. We see text in terms of words, not in terms of letters, and the meanings of those words must be in long-term

[3]Of course in this example, it is not simply the chunking that makes the numbers easy to recall. It is the strong linkage with other, semantic knowledge.

storage. The third point is something we will encounter again and again in different contexts: How students respond to a piece of information presented to them depends both on what they know already and on their mental state—what information they are cued to access.

The number of chunks a piece of information has for an individual depends not only on whether or not they have relevant associations but on how strong and easily accessible that knowledge is in long-term memory. When a bit of knowledge—a fact or process—is easily available and can easily be used as a single unit in working memory, we say the knowledge is *compiled.* Computer programming is a reasonably good metaphor for this. When code in a high-level computer language has to be translated line by line into machine instructions, the code runs slowly. If the code is compiled directly so that only machine-language instructions are presented, the code runs much more quickly.

Some of the difficulties students encounter—and that we encounter in understanding their difficulties—arise from this situation. Physics instructors work with many large blocks of compiled knowledge. As a result, many arguments that seem simple to them go beyond the bounds of working memory for their students. If the students have not compiled the knowledge, an argument that the instructor can do in a few operations in working memory may require the student to carry out a long series of manipulations, putting some intermediate information out to temporary storage in order to carry out other parts of the reasoning.

Studies with subjects trying to recall strings of information indicate that items fade from working memory in a few seconds if the subject does not try to remember the information by repeating it consciously [Ellis 1993]. This working memory repetition is known as *rehearsal.* Think about looking up a telephone number in a phonebook. Most of us can't remember it—even for the few seconds needed to tap in the number—without actively repeating it.

The short lifetime of working memory has serious implications for the way we communicate with other people, both in speaking and in writing. In computer science, the holding of information aside in preparation for using it later is called *buffering*, and the storage space in which the information is placed is called a *buffer.* Since human memory buffers are volatile and have a lifetime of only a few seconds, it can be very confusing to present information that relies on information that has not yet been provided. Doing this can interfere with a student's ability to make sense out of a lecture or with a websurfer's ability to understand a webpage.[4]

Working memory contains distinct verbal and visual parts. A second characteristic of working memory that has been well documented in the cognitive literature[5] is that working memory contains distinct components. At least the verbal component (the *phonological loop*) and the visual component (the *visual sketchpad*) of working memory appear to be distinct. (I am not aware of any evidence in the cognitive literature about the independence of other components such as mathematical or musical.) This has been demonstrated by showing that two verbal tasks or two spatial tasks interfere with each other substantially more than a visual task interferes with a verbal task, or vice versa.[6]

[4] In the theory of communications, this leads to the "given-new principle" in conversation [Clark 1975] and writing [RedishJ 1993].

[5] See [Baddeley 1998] and [Smith 1999] and the references therein.

[6] The evidence for this is also strong from neurophysiological lesion studies [Shallice 1988].

2. Long-term memory

Long-term memory is involved in essentially all of our cognitive experiences—everything from recognizing familiar objects to making sense of what we read. An important result is that recall from long-term memory is *productive* and *context dependent*.

Long-term memory is productive. What we mean here by "productive" is that memory response is active. Information is brought out of long-term storage into working memory and processed. In most cases, the response is not to simply find a match to an existing bit of data, but to build a response by using stored information in new and productive ways. This construction is an active, but in most cases, an automatic and unconscious process. Think of language learning by a small child as a prototypical example. Children create their own grammars from what they hear.[7] Another model of the recall process is computer code. A result, such as sin(0.23 rad), may be stored as tables of data from which one can interpolate or as strings of code that upon execution will produce the appropriate data. Analogs of both methods appear to be used in the brain.

Another example demonstrates that it's not just sensory data that the brain is processing; cognitive processes such as recall and identification of objects is also productive. Look at the picture in Figure 2.1. It consists of a number of black blobs on a white background. The subject of the picture will be immediately obvious to some, more difficult to see for others. (See the footnote if you have looked at the picture for a while and can't make it out.[8]) Even though you may never have seen the particular photograph from which this picture was constructed, your mind creates recognition by pulling together the loosely related spots and "constructing" the image. Once you have seen it, it will be hard to remember what it looked like

Figure 2.1 A picture of an animal [Frisby 1980].

[7] The fact that they don't always create the same rules as their parents have is one of the facts that causes languages to evolve.

[8] The picture is of a Dalmatian dog with head down and to the left, drinking from a puddle in the road, seen through the shadows under a leafy tree.

to you when you couldn't see it. When you couldn't "see" the picture in the blobs, was there a picture there? Now that you see it, where is the picture, on the paper or in your brain?

Long-term memory is context dependent. By *context-dependent*, I mean that the cognitive response to a mental stimulus depends on both (1) the external situation and the way in which the stimulus is presented and (2) the state of the respondent's mind when the stimulus is presented. The first point means, for example, that for a problem on projectile motion presented to a student in a physics class, a student might bring out of long-term memory a different repertoire of tools than the ones she might access if the same problem arose on the softball field. To see what the second point means, consider for example, a situation in which a student is asked to solve a physics problem that could be solved using either energy or force methods. If the problem is preceded by a question about forces, the student is much more likely to respond to the problem using forces than if the question were not asked [Sabella 1999].

To show how the context dependence affects the resources that one brings to the analysis of a situation, let's look at the following example.[9] Suppose I am holding a deck of 3×5 file cards. Each card has a letter on one side and a number on the other. I remove four cards from the deck and lay them on the table as shown in Figure 2.2.

I make the following claim*: This set of four cards satisfies the property that if there is a vowel on one side of the card, then there is an odd number on the other.* How many cards do you need to turn over to be absolutely certain that the cards have been correctly chosen to satisfy this property?

Try to solve this problem before looking to the footnote for the answer.[10] Careful! You have to read both the claim and the question carefully. A similar problem but with a different context is the following.

You are serving as the chaperone and bouncer[11] at a local student bar and coffee house. Rather than standing at the door checking IDs all the time, you have occupied a table so you can do some work. When patrons come in and give their order, the servers bring you cards

Figure 2.2 An abstract problem. (See text.)

[9] Adapted from [Wason 1966] and [Dennett 1995].

[10] You may have to turn over at most two cards to be sure I am telling the truth. (If the first card fails, you know I am wrong right away.) The only cards that are relevant are the number "2" and the letter "A." Note that the statement only says "if," not "if and only if." To test whether $p \rightarrow q$, you have to test the equivalent statements: $p \rightarrow q$ and $\sim q \rightarrow \sim p$.

[11] This is a highly culture-dependent example. In order to solve it, you must know that in most American communities, the law prohibits the purchase of alcoholic beverages by individuals younger than 21 years of age. Putting the problem into this cultural context also broadens the concerns of some respondents as well as their tools for solving it. Some worry whether or not the servers can be trusted or might be lying if, for example, a friend were involved.

Figure 2.3 A more concrete card problem. (See text.)

with the patron's order on one side and their best guess of the patron's age on the other. You then decide whether to go and check IDs. (The servers can be assumed to be trustworthy and are pretty good guessers.)

A server drops four cards on the table. They land as shown in Figure 2.3. Which cards would you turn over in order to decide whether to go back to the table to check IDs?[12]

This problem is mathematically isomorphic to the previous one, yet most American adults find the first problem extremely difficult. They reason incorrectly or read the problem in an alternative way that is plausible if a word or two is ignored or misinterpreted. I have presented these problems in lectures to physics departments many times. More than two-thirds of physicists produce the wrong answer to the K2A7 problem after a minute or two of consideration. Almost everyone is able to solve the second problem instantly.

These problems provide a very nice example of both productive reasoning and context dependence. In the two cases, most people call on different kinds of reasoning to answer the two problems. The second relies on matching with social experience—a kind of knowledge handled in a much different way from mathematical reasoning.

This result has powerful implications for our attempts to instruct untrained students in physics. First, it demonstrates that the assumption that "once a student has learned something, they'll have it" is not correct. Most physics instructors who have tried to use results the students are supposed to have learned in a math class are aware of this problem. The example shows that even changing the context of a problem may make it much more difficult. Second, it points out that a problem or reasoning that has become sufficiently familiar to *us* to feel like 16/Coke/52/Gin & tonic may feel like K2A7 to our students! We need to maintain our patience and sympathy when our students can't see a line of reasoning that appears trivial to us.

Long-term memory is structured and associative. The examples in the previous subsections illustrate the fundamental principle that the activation of knowledge in long-term memory is structured and associative. When a stimulus is presented, a variety of elements of knowledge may be activated (made accessible to working memory). The particular elements that are activated can depend on the way the stimulus is presented and on the state of the mental system at the time (the context). Each activation may lead to another in a chain of *spreading activation* [Anderson 1999].

[12] You would only have to turn over the cards labeled "16" and "Gin & Tonic." You are not concerned with what a person clearly much older than 21 orders, and anyone is allowed to order a coke.

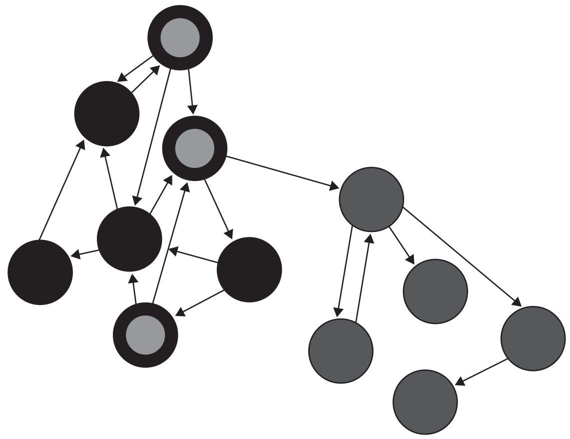

Figure 2.4 An example illustrating linked context-dependent schema.

The key to understanding student reasoning is understanding the *patterns of association* that activate knowledge elements. In general, a pattern of association of knowledge elements is sometimes referred to as a *knowledge structure.* A pattern that tends to activate together with a high probability in a variety of contexts is often referred to as a *schema* (plural, *schemas* or *schemata*) [Bartlett 1932] [Rumelhart 1975]. This is illustrated in Figure 2.4. Each circle represents an element of knowledge. Arrows indicate a probability that the activation of one element will lead to the activation of another. The different colorings of the circles indicate schemas—associated knowledge elements that tend to be activated together. Notice that on the left, some elements have multiple circles. This indicates that particular knowledge elements can activate different schemas, depending on context.

As an example, consider meeting a new person at a beach party. In your conversation with this individual, you activate a number of responses—seeking knowledge of your own about topics the other person raises, looking for body language that demonstrates interest in continuing the conversation, and so on. If, at some point later in the party, the individual falls and is knocked unconscious, a different chain of knowledge and responses is activated. Is the person seriously injured? Do you need to get them to an emergency room? Should the person be moved or medical personnel called? You still begin with a response to the individual as a person, but different associated knowledge elements are activated.[13]

When a schema is robust and reasonably coherent I describe it with the term *mental model.* Since scientific models tend to be organized around the existence, properties, and interaction of objects, when a mental model has this character I refer to it as a *physical model.* A physical model may or may not agree with our current community consensus view of physics. For example, the phlogiston picture of thermodynamics was organized around an

[13] In reading this situation, did you envision the person as an individual as the same or opposite sex to your own? This component of context is an example of context affecting activation responses.

image of a physical substance that we now understand cannot be made consistent with physical observations, so this physical model does not agree with our current community physical model. Two individuals well versed in a topic may have different physical models of the same system. For example, a circuit designer may model electric circuits in terms of resistors, inductors, capacitors, and power sources—macroscopic objects with particular properties and behaviors. A condensed-matter physicist may think of the same system as emerging from a microscopic model of electrons and ions.

Cognitive resources for learning

The fact that the mind works by context-dependent patterns of association suggests that students reason about physics problems using what they think they know by generalizing their personal experience. This doesn't sound surprising at first, but we may be surprised at some of its implications.

When I learned[14] that students in introductory physics often bring in a schema of motion that says objects tend to stop, I enthusiastically presented the result to Sagredo. He was skeptical and argued "If you don't tell them about friction, they wont know about it." Sorry, Sagredo. They may not know the word "friction" (they mostly do) or the rules of physics that describe it (they mostly don't), but they are very familiar with the fact that if you push a heavy box across the floor, it stops almost as soon as you stop pushing it. They also know that if they want to move, they have to exert an effort to walk and when they stop making that effort, they stop.

"But," responds Sagredo, "if you push a box very hard along a slippery floor it will keep going for quite a distance. If you run and stop making an effort, you'll continue going and fall over. Surely they know those facts as well." Absolutely right, Sagredo. But the problem is that most students do not attempt to make a single coherent picture that describes all phenomena. Most are satisfied with a fairly general set of often-inconsistent rules about "how things work."

Over the years there have been some disagreements among researchers as to the nature of the schemas students bring with them to the study of introductory physics. Some researchers suggested that students had "alternative theories"—reasonably self-consistent models of the world that were different from the scientist's [McCloskey 1983]. But extensive work by a number of different researchers (see in particular the work of McDermott [McDermott 1991] and diSessa [diSessa 1993]) suggests that student knowledge structures about physics tend to be weak patterns of association that rarely have the character of a strong schema.[15]

Although I concur with this view, I note that occasionally a student's schema may be more coherent than we as scientists tend to give credit for, since we analyze the student's views through the filter of our own schemas. For example, a student who thinks about "motion" and fails to separate velocity from acceleration may seem inconsistent to a physicist, whereas, in fact, the student may feel he is consistent but may have a model that only works for a rather limited set of specific situations and questions. As long as the student is either aware

[14] I learned this from reading the paper of Halloun and Hestenes [Halloun 1985a] that also inspired Mazur—see chapter 1.

[15] By "strong" or "weak" we simply mean a high or low probability of activating a link to other related (and appropriate) items in a student's schema.

of those limitations or not presented with situations in which that model doesn't work, the student can function satisfactorily.

There are two reasons why it is important for us to understand the knowledge and reasoning about the physical world that students bring with them into our classes. First, it helps us understand what errors the students will commonly make and how they will misinterpret what we say and what they read. We can then use this understanding to help us design both new and improved instruction and better evaluations. I find understanding student thinking particularly useful in helping me to answer students' questions, both one-on-one during office hours and during lecture. It is very easy to misinterpret a student's question as being much more sophisticated than it really is. Knowing common lines of reasoning or associations can help probe the individual student as to what his or her problem really is. In lecture it can help us understand what a question that seems a "stupid" question on its face may really be telling us about the state of many of the students in our class and our interaction with them.

Second, the students bring into our class the basic *resources* from which they will build their future knowledge. Since new knowledge is built only by extending and modifying existing schema, students' existing knowledge is the raw material we have to work with to help them build a more correct and more scientific knowledge structure [Hammer 2000] [Elby 1999].

The knowledge and reasoning that students bring into our class have been analyzed in three ways that are useful for us: (1) as common naïve conceptions, (2) in terms of primitive reasoning elements, and (3) in terms of the way reasoning and knowledge are situated in everyday real-life circumstances. The last-named is called *situated cognition.*

1. Robust reasoning structures: Common naïve conceptions

Occasionally, students' patterns of associations concerning physical phenomena are strikingly robust—they occur with a high enough probability in many contexts for us to refer to them as mental models. In many cases, they contain inappropriate generalizations, conflations of distinct concepts (such as treating velocity and acceleration as a single concept, "motion"), or separations of situations that should be treated uniformly (such as treating a box sliding along a rough floor and a rapidly moving baseball using different rules). When a particular mental model or line of reasoning is robust and found in a significant fraction (say on the order of 20% of students or more), I refer to it as a *common naïve conception.* In the education research literature, these patterns are often referred to as *misconceptions*, *alternative conceptions*, or *preconceptions*, particularly when they lead to incorrect predictions or conclusions. I choose the more descriptive term, "common naïve conceptions," rather than the most common parlance, "misconceptions," both because it lacks the pejorative sting and because I want to emphasize the complexity of the student concept. Usually, these conceptions are not "just wrong." Students may be naïve, but they're not fools. Their naïve conceptions are usually valuable and effective in getting them through their daily lives. Indeed, most naïve conceptions have kernels of truth that can help students build more scientific and productive concepts.

The presence of common naïve conceptions really isn't so surprising if we think about our students' previous experience. Why should we be surprised that students think that any moving object will eventually come to a stop? In their direct personal experience that is always the case. It's even the case in the demonstrations we show in class to demonstrate the opposite! When I slide a dry-ice levitated puck across the lecture table, I catch it and stop it at the end of the table. If I didn't, it would slide off the table, bounce, roll a short distance, and stop. Every student knows that. Yet I ask them to focus on a small piece of the demonstration—the stretch of about four or five seconds when the puck is sliding along the table smoothly—and extend that observation in their minds to infinity. The student and the teacher may focus on different aspects of the same physical phenomena.[16]

Many teachers show surprise when they learn the results of physics education research demonstrating that students regularly generalize their naïve schemas incorrectly. Why should it be surprising that students think cutting off part of a lens will result in only part of an image being visible on the screen [Goldberg 1986]? Try looking through a magnifying glass! (Yes, I know that's not a real image.) Where do students get the idea that electricity is something that is used up in a resistor [Cohen 1983]? We've told them that you need circuits and that currents flow in loops! Although we don't always think about it, most of our students have had extensive experience with electricity by the time they arrive in our classes. When I said the current had to come in one wire and go out the other, one of my students complained: "If all the electricity goes back into the wall, what are we paying for?"

Much of the effort in the published physics education research literature has been to document the common naïve conceptions of introductory physics students and to develop instructional methods to deal with them. To get an introduction to this literature, consult the papers listed in the Resource Letter given on the Resource CD [McDermott 1999]. A good overview is given in the books [Arons 1990], [Viennot 2001], and [Knight 2002].

2. Modular reasoning structures: Primitives and facets

Perhaps the most extensive and detailed analysis of student reasoning in introductory physics has been diSessa's monumental work, "Toward an Epistemology of Physics" [diSessa 1993].[17] In this study, diSessa analyzes the evolution of reasoning in mechanics of 20 MIT students in introductory calculus-based physics. Although this is a fairly small number of students and a rather narrow population, the care and depth of the analysis make it worthy of attention.[18] Subsequent investigations show the presence of diSessa's results in much broader populations. As of this writing, diSessa's approach has only rarely been applied to the development of new curriculum. Nonetheless, because of the powerful insights it provides into student reasoning, I believe it will be of use both in future curriculum reform and in research trying to understand student thinking, and so I include a brief discussion of his ideas here.

DiSessa investigated people's *sense of physical mechanism*, that is, their understanding of "Why do things work the way they do?" What he found was that many students, even after instruction in physics, often come up with simple statements that describe the way they think things function in the real world. They often consider these statements to be "irreducible"—as the obvious or ultimate answer; that is, they can't give a "why" beyond it. "That's just the way things work," is a typical response. DiSessa refers to such statements as *phenomenological primitives.*

Some of these primitives may activate others with a reasonably high priority, but diSessa claims that most students have rather simple schemas. Primitives tend to be linked directly to a physical situation. They are recognized in a physical system rather than derived by a long chain of reasoning.

[16] This argument is made in a slightly different context in [KuhnT 1970].

[17] A shorter, more accessible introduction to diSessa's ideas is given in [diSessa 1983].

[18] Note that the point of this kind of study is to determine the range of possibilities, not their distribution among a particular population. As a result, the small number of students in the study is not a serious drawback.

As an example, consider the response of a student to pushing a box on a rough surface. The student might respond that "you need a big force to get it going" (*overcoming*: one influence may overpower another by increasing its magnitude), but then "you need a force to keep it going" (*continuous push*: a constant effort is needed to maintain motion). DiSessa identifies the parentheticals as primitives. Notice that the primitives are neither wrong nor right in themselves. They are certainly correct in some circumstances, and diSessa points out that experts use many primitives as quick and easy bits of compiled knowledge—but they are linked so as to only be used in appropriate circumstances.

I like to add an additional structure. Some of diSessa's phenomenological primitives are very abstract (*Ohm's primitive*, for example), and others refer to reasonably specific physical situations (*force as mover*, for example). I prefer to distinguish abstract reasoning primitives from those primitives applied in a particular context, which I refer to as *facets* (following Minstrell).[19] What I call an *abstract reasoning primitive* has a general logical structure, such as "if two quantities *x* and *y* are positively related, more *x* implies more *y*." What I call a facet implies a *mapping* of the slots in the primitive into particular variables in a particular physical context. This is illustrated in Figure 2.5. As diSessa points out [diSessa 1993], there are a very large number of facets, corresponding to the complexity of living in the world. In my formulation, this complexity is seen as arising from mapping a reasonably small number of abstract reasoning primitives (perhaps a few dozen) onto the large diversity of physical situations.

An example of mapping might be, "if the liquid is higher in one of the glasses, there is more liquid in that glass." In one of Piaget's classic experiments, children are shown a vessel containing some water. The water is then poured into a narrower vessel so that it rises to a higher level. Before the age of about five years, most children say that the amount of water has increased (because it's higher) or decreased (because it's narrower).[20] Both those children

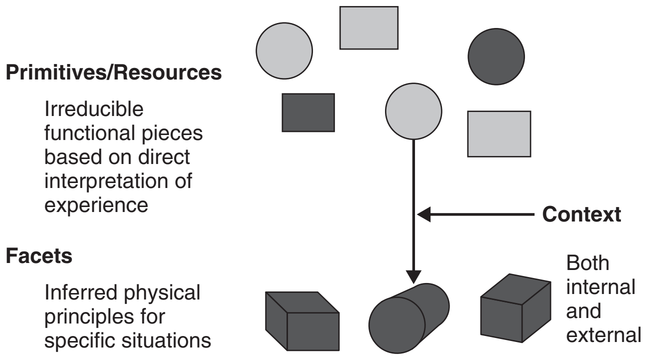

Figure 2.5 A visual representation of the mapping of abstract primitives into specific facets in a particular physical context.

[19] The term *facet* was introduced by Jim Minstrell [Minstrell 1992]. Minstrell listed large numbers of reasoning statements, some correct, others not, that students gave as explanations for their observations or in predictions about specific physical contexts.

[20] This is not a simple failure to understand the meaning of the question. A child may get upset that a sibling is "getting more" even if she is shown that when poured into a glass similar to their own they have the same heights.

who see it as more and those who see it as less are using what is essentially the same abstract reasoning primitive, but with different mappings (focusing on height or width). After the age of about five, children learn a "compensation" abstract reasoning primitive, something like "if two effects act on a variable, one to increase it and the other to decrease it, those effects tend to cancel." Children reaching this ability to reason are said to have achieved *Piagetian conservation.*

I've seen something very similar to this with engineers in first-year calculus-based physics at the University of Maryland. We were discussing the collision of two objects of different masses, and I asked about the relative size of the forces each object experiences. One group of students said, "The larger objects feel a bigger force since they have bigger masses." A second group said, "The smaller objects feel a bigger force since they change their velocity more." Only a small number of students had reached the "Piagetian conservation" stage of activating the compensation implicit in Newton's second law.

This kind of approach—analyzing the responses of our students in terms of primitives and facets—helps us understand more clearly the kinds of reasoning we can expect. The critical realization that arises from this kind of analysis is that students' common naïve conceptions are not simply wrong. They are based on correct observations but may be generalized inappropriately or mapped onto incorrect variables. If we can extract elements that are correct from students' common reasoning, we can build on these elements to help students reorganize their existing knowledge into a more complete and correct structure.

3. Activating resources from everyday experience: Situated cognition

The primitives discussed above tend to refer to specific real-world situations as asked about or observed in a physics class. A group of education specialists have focused on the difference between day-to-day reasoning and the kind of reasoning taught and learned in schools.

Sagredo once stopped me in the hall after his introductory physics class for physics majors. "You'll never guess what they couldn't do today! I was talking about projectile motion and asked them to describe what happens when a kicker kicks a football. I just wanted a description of the process—the ball goes up, travels a ways down the field, and comes down. No one would say *anything*, even when I pressed them. Why couldn't they give me a simple description?" I suspect, Sagredo, that it was because they weren't really sure what you wanted. They might well have expected that you wanted some technical or mathematical description in terms of forces, graphs, velocities, and accelerations. If they just said what you really wanted—the simple day-to-day physical description of the process—they were afraid they would look foolish. They may have been right to respond that way, given their previous experience with physics classes.

Most instruction in the United States today, despite reform efforts, continues to bear little relation to students' everyday lives. But many of the skills we are trying to teach can be tied to reasoning skills that the students possess and use every day. An interesting example comes from middle school arithmetic [Ma 1999]. Consider the following problem.

> A group of students has 3½ small pizzas, each whole divided into 4 parts. How many students can have a piece?

The reasoning used by a student to solve this problem is quite a bit different from the algorithm one learns for dividing fractions ($3\frac{1}{2}/\frac{1}{4} = 4 \times 3\frac{1}{2}$). The student might say something like: "Each pie can serve 4 students, so the 3 pies can serve 12. The ½ pie can serve 2, so a total of 14 can have a piece." This reasoning, like the reasoning we use to solve the Coke/Gin & Tonic problem in the previous section, relies on nonformal thinking that is linked to our everyday social experience. Tying it to the dividing by fractions problem—showing that division means finding how many times a part can be found in the whole—can help students make sense of what division really means and why division is a useful concept.

The use of context knowledge to help solve problems is a common feature of how people reason their way through situations in everyday life. A group of educators led by Jean Lave and Lucy Suchman [Lave 1991] [Suchman 1987] places this cognitive fact at the center of their educational reform efforts, creating *cognitive apprenticeships* and using *situated cognition*. There is an extensive educational literature on this topic,[21] and some dramatic improvements have been gained in children's understanding of and effectiveness using arithmetic by finding ways to use these everyday resources.

IMPLICATIONS OF THE COGNITIVE MODEL FOR INSTRUCTION: FIVE FOOTHOLD PRINCIPLES

Any model of thinking is necessarily complex. We think about many things in many ways. In order to find ways to see the relevance of these cognitive ideas and to apply them in the context of physics teaching, I have selected five general principles that help us understand what happens in the physics classroom.

1. The constructivism principle
2. The context principle
3. The change principle
4. The individuality principle
5. The social learning principle

1. The constructivism principle

> *Principle 1:* Individuals build their knowledge by making connections to existing knowledge; they use this knowledge by productively creating a response to the information they receive.

This principle summarizes the core of the fundamental ideas about the structure of long-term memory and recall. The basic mechanism of the cognitive response is context-dependent association. A number of interesting corollaries, elaborations, and implications that are relevant for understanding physics teaching come from the constructivism principle.[22]

[21] See, for example, [Lave 1991]. A very readable introduction to the subject is [Brown 1989].

[22] These properties point out that our whole structure of patterns of association/schemas/mental models is a somewhat fuzzy one. The boundaries between the different structures are not sharply delineated. There are lots of examples of this sort of description in physics. Consider, for example, the excitations of a crystal lattice. We can describe the excitations in terms of phonons or in terms of continuous waves. In some limiting cases, it is clear which description is the more useful; in others, they may overlap.

Some of the characteristics of schemas clarify what is happening when students make mistakes. Often in listening to my students explain what they think, I used to become confused and sometimes irritated. How can they say *x* when they know the contradictory principle *y*? Why can't they get started on a problem when they certainly know the relevant principle? They just told it to me two minutes ago! Why did they bring up that particular principle now? It doesn't apply here! The well-documented characteristics of mental structures described in the first part of this chapter help us understand that these sorts of errors are natural and to be expected.

It also makes us realize that we must get a lot more feedback than we traditionally get if we want to understand what our students are really learning. Traditional testing often fails to show us what students really think or know because many different schemas can produce the correct solution to a problem. Even if a student goes through the same steps as we do, there's no guarantee that their schema for choosing the steps is the same as ours.[23] I once asked a student, who had done a homework problem correctly, to explain his solution. He replied: "Well, we've used all of the other formulas at the end of the chapter except this one, and the unknown starts with the same letter as is in that formula, so that must be the one to use."

Part of the way we fool ourselves with standard testing methods is that we are interested "in what students know." If they don't access the right information in an exam, we give them clues and hints in the wording to activate access. But since essential components of a schema are the processes for access to information, we are not testing the students' patterns of associations if we narrow the context and provide detailed cues. The student "has" the information, but it is inert and cannot be used or recalled except in very narrow, almost preprogrammed situations.

To find out what our students really know, we have to find out what resources they are bringing and what resources they are using. We have to give them the opportunity to explain what they are thinking in words. We must also only give exam credit for reasoning and not give partial credit when a student tries to hit the target with a blast of shotgun pellets and accidentally has a correct and relevant equation among a mass of irrelevancies. To know whether our students access the information in appropriate circumstances, we have to give them more realistic problems—problems that relate directly to their real-world experience and do not provide too many "physics clues" that specify an access path for them. (I'll discuss the implications of this for assessment in chapter 4.)

2. The context principle

The second principle reminds us of the nonuniqueness of the cognitive response and sets the stage for the description of the dynamics of building mental structures.

> *Principle 2:* What people construct depends on the context—including their mental states.

It's very easy to forget the warnings and drop back into the model that assumes students either know something or don't. Focusing on the context dependence of a response helps us

[23] The difficulty is that the mapping from underlying schema to problem-solving steps is not one-to-one. A specific example of this is given in [Bowden 1992].

keep in mind that the situation is not that simple. Nice examples of context dependence in students' responses to physics abound (although they are sometimes not presented that way). One particularly clear example comes from the work of Steinberg and Sabella [Steinberg 1997] [Sabella 1999]. At the end of the first semester of engineering (calculus-based) physics at the University of Maryland, they gave a section of 28 students the (to an expert) equivalent pair of questions, shown in Figure 2.6.

The first question is taken from the Force Concept Inventory (FCI) [Hestenes 1992].[24] It is stated in common speech using familiar everyday objects. The second is a typical physics class situation. It is abstract, and it involves idealized laboratory-style objects. The FCI was given as an ungraded diagnostic at the end of classes[25] and the problem as part of the final exam one week later. Both test the students' understanding of Newton's first law.

Although 25 (~90%) got the exam question correct, only 15 of the students (~55%) got the FCI question right. Nearly half of the students who succeeded with the problem in the exam context missed the question presented in a nonphysics context (11/25 ~ 45%). Interviews with the students suggest that this is a real phenomenon, not just a result of an additional week's time to study.

An elevator (as illustrated) is being lifted up an elevator shaft by a steel cable. When the elevator is moving up the shaft at a constant velocity (assume that any frictional forces due to air resistance can be ignored):

a) the upward force on the elevator by the cable is greater than the downward force of gravity

b) the amount of upward force on the elevator by the cable is equal to the downward force of gravity

c) the upward force on the elevator by the cable is less than the downward force of gravity

d) it goes up because the cable is being shortened, not because of the force being exerted by the elevator on the cable

e) the upward force by the elevator on the cable is greater than the downward force due to the combined effects of air pressure and the force of gravity.

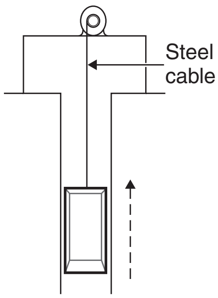

Ignore all friction and air resistance in this problem.

A steel ball resting on a small platform mounted to a hydraulic lift is being lowered at a constant speed, as shown in the figure at right.

a) Draw a free-body diagram of the ball. Describe each type of force.

b) Compare the magnitudes of the forces you have drawn. Explain your reasoning.

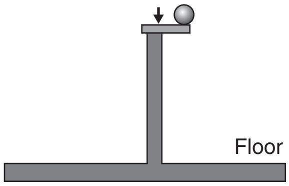

Figure 2.6 Two problems from [Steinberg 1997] demonstrating the context dependence of student responses.

[24] The FCI is discussed in detail in chapter 5 and is included in the Resource CD associated with this volume.

[25] These tests are often given as ungraded in order to encourage the students to give the answer they believe rather than the answer they might think we want in physics class.

What we really want to help our students do is build their knowledge into a coherent schema that is appropriately linked and that is triggered in a wide range of appropriate contexts.

3. The change principle

This principle deals with the dynamics of the mental state. It states that schemas are not only the way that we organize our interactions with the world, but they also control how we incorporate new information and experiences [Bransford 1973].

> *Principle 3:* It is reasonably easy to learn something that matches or extends an existing schema, but changing a well-established schema substantially is difficult.

I pose a number of restatements and elaborations of this principle as corollaries to clarify what it means for teaching.

> *Corollary 3.1:* It's hard to learn something we don't almost already know.

All students have things they know (some of which may be wrong!), things they are a bit familiar with, and things they have no knowledge about at all.

I like to look at this as an archery target. What they know is the bull's-eye—a compact black area; what they know a little about is a gray area surrounding the black; outside that is a white "rest of the world" about which they are clueless. To teach them something, we do best to hit in the gray. A class full of students is a challenge because all of their gray areas are not the same. I want to hit as many of the grays as possible with each paint-tipped shaft of information to turn gray areas black.[26] This metaphor only describes some aspects of the process. The real issue is that when we "hit in the gray," the student has many appropriate links that can be made to weave the new knowledge into their existing structure in an appropriate way.

In communication studies, an important implication of this corollary is called the *given-new principle* [Clark 1975] [RedishJ 1993]. It states that new information should always be presented in a context that is familiar to the reader and that the context should be established first. The analogous statement is very important in physics teaching, especially at the introductory level. As physicists with years of training and experience, we have a great deal of "context" that our students don't possess. Often we are as fish in water; unaware of having this context and unaware that it is missing in our students.

We can cite a number of specifics that are violations of the given-new principle. One important example is that we often use terms that students are not familiar with—or that they use in a different sense than we do. Lakoff and Johnson [Lakoff 1980], as a part of their study of the way speakers of English build their meaning of the term *force,* classified the characteristics of common metaphors using the term. Among their list of 11 characteristics, 8 involved

[26] This picture also interacts strongly with the social learning principle discussed below. Items in the gray are those that the student can learn via social interactions with teachers or more expert students. The followers of the Russian psychologist Lev Vygotsky give the gray region the unfortunate name "zone of proximal development" [Vygotsky 1978].

the will or intent of an individual! But most of us are so familiar with the technical meaning of "force" that we are surprised to learn that a significant fraction of our introductory students do not believe that a table exerts a force on a book it is supporting [Minstrell 1982]. Why doesn't the book fall through? The table is just "in the way." (This issue is discussed in more detail under the heading "Bridging" later in this chapter.)

The problem caused by the interpretation of common speech words for technical ones is not simple. I know that the terms "heat" and "temperature" are not really distinguished in common speech and are used interchangeably for the technical terms "temperature" (average energy per degree of freedom), "internal energy", and "heat flow" (flow of internal energy from one object to another). In one class, I stated this problem up front and warned my students that I would use the terms technically in the lecture. Part way through I stopped, realizing that I had used the word "heat" twice in a sentence—once in the technical sense, once in the common speech sense.[27] It's like using the same symbol to stand for two different meanings in a single equation. You can occasionally get away with it,[28] but it really isn't a good idea!

Putting new material in context is only part of the story. Our students also have to see the new material as having a plausible structure in terms of structures they know. We can state this as another useful corollary.

Corollary 3.2: Much of our learning is done by analogy.

This and the previous corollary make what students know at each stage critical for what we can teach them. Students, like everyone else, always construct their knowledge, and what they construct depends on how what we give them interacts with what they already have. This has important implications for the development of instructional techniques that help students overcome strong misconceptions. (See the discussion of "bridging" in a later section.)

One implication of these results is that we should focus on building structures that are useful for our students' future learning. I state this as a third corollary.

Corollary 3.3: "Touchstone" problems and examples are very important.

By a *touchstone problem*,[29] I mean one that the student will come back to over and over again in later training. Touchstone problems often become the analogs on which they will build the more sophisticated elements of their schemas. It becomes extremely important for students to develop a collection of a few critical things that they really understand well.[30] These become the "queen bees" for new swarms of understanding to be built around. I

[27] "If there is no heat flow permitted to the object, we can still heat it up by doing work on it."

[28] I have seen colleagues write the energy levels of hydrogen in a magnetic field as $E_{nlm} = E_n - \left(\frac{e\hbar}{2m}\right)mB$ where the m in the denominator is the electron mass and the one in the numerator is the z-component of the angular momentum. Most physicists can correctly interpret this abomination without difficulty.

[29] In his discussion of scientific paradigms, T. S. Kuhn refers to these problems as *exemplars* [KuhnT 1970].

[30] In addition to giving them centers on which to build future learning, knowing a few things well gives the student a model of what it means to understand something in physics. This valuable point that has been frequently stressed by Arnold Arons [Arons 1990]. It is an essential element in developing scientific schemas.

believe the sense that some superficially uninteresting problems serve this role is the reason they have immense persistence in the community. Inclined plane problems really aren't very interesting, yet the occasional suggestions that they be done away with are always resisted vigorously. I think the resisters are expressing the (often unarticulated) feeling that these are the critical touchstone problems for building students' understanding of vector analysis in the plane.

Corollary 3.3 is one reason we spend so much time studying the mass on a spring. Springs are of some limited interest in themselves, but small-amplitude vibrations are of great general importance. The spring serves as a touchstone problem for all kinds of harmonic oscillations from electrical circuits up to quantum field theory.

Analyzing a curriculum from the point of view of the schema we want students to develop, their preexisting schemas, and touchstone problems that will help them in the future can help us understand what is critical in the curriculum, which proposed modifications could be severely detrimental, and which might be of great benefit.

Combining these ideas with the idea of associations discussed under Principle 1 leads us to focus on the presence of a framework or structure within a course. It suggests that building a course around a linked series of touchstone problems could be of considerable assistance in helping students understand the importance and relevance of each element. Such a structure is sometimes referred to as a *story line*.

Unfortunately, if students are *not* blank slates, sometimes what is written is—if not entirely wrong—inappropriate for future learning in physics.[31] Then it can seem as if we have run into a brick wall. This brings us to the next corollary.

Corollary: 3.4: It is very difficult to change an established mental model.

Traditionally, we've relied on an oversimplified view of Principle 1, the constructivism principle, to say: "Just let students do enough problems and they'll get the idea eventually." Unfortunately, this simple translation of the principle doesn't necessarily work. Although practice is certainly necessary to help students compile skills into easily retrievable knowledge, there is no guarantee that they will link those skills into a structure that helps them to understand what's going on and how to use the basic concepts appropriately.

The limitations of doing lots of problems were investigated in a study done in Korea. Eunsook Kim and Jae Park looked at the response of 27 students in an introductory college physics class to the FCI [Kim 2002]. American students at large, moderately selective state universities (such as the University of Maryland, the Ohio State University, or the University of Minnesota) who have taken one year of high school physics score an average of 45–50% on this test before beginning a calculus-based physics class. The students Kim studied had taken an apparently much more rigorous high school physics program in which each student had done an average of about 1500 problems (ranging between 300 and 2900) end-of-chapter problems. In a typical American high school class, students will do 300 to 400 such problems. Despite doing 5 to 10 times as many problems as American students, the students

[31] Perhaps a *palimpsest* is a better metaphor for a student's mind than a blank slate. According to the American Heritage Dictionary, a palimpsest is "a manuscript, typically of papyrus or parchment, that has been written on more than once, with the earlier writing incompletely erased and often legible."

still had substantial conceptual difficulties with fundamental concepts of mechanics at rates comparable to those seen with American students. There was little correlation between the number of problems the students had done and the conceptual understanding displayed.

This study and others like it are a warning against relying on the idea that "repetition implies effective learning"—that is, that frequent repetition of a particular type of activity is *all* that is needed to produce a useful functional schema. Repetition *is* necessary to create compiled knowledge, but it is not sufficient. For effective usage, the compiled element needs to be linked into a coherent schema about the subject.

Once students learn how to do problems of a particular type, many will learn nothing more from doing more of them: New problems are done automatically without thinking. This also means that testing by varying homework problems slightly may be inadequate to probe the student's schemas. More challenging tests involving a variety of modes of thinking (problem solving, writing, interpreting, organizing) are required. Such testing is discussed in detail in chapter 4.

It has been demonstrated over and over again that simply telling students some physics principle doesn't easily change their deep ideas.[32] Rather, what often happens is that instead of changing their schema substantially, a poorly linked element is added with a rule for using it only in physics problems or for tests in one particular class. This and the fact that a schema can contain contradictory elements is one possible reason "giving more problems" is often ineffective.

A few years ago, I heard a lovely anecdote illustrating the barriers one encounters in trying to change a well-established mental model.[33] A college physics teacher asked a class of beginning students whether heavy objects fall faster than light ones or whether they fall at the same rate. One student waved her hand saying, "I know, I know." When called on to explain she said: "Heavy objects fall faster than light ones. We know this because Galileo dropped a penny and a feather from the top of the leaning tower of Pisa and the penny hit first." This is a touchstone example for me. It shows clearly that the student had been told—and had listened to—both the Galileo and the Newton stories. But she had transformed them both to agree with her existing mental model.[34]

Fortunately, mechanisms are available for helping students restructure even well-established mental models; these methods are discussed later in this chapter.

4. The individuality principle

One might be tempted to say: Fine. Let's figure out what the students know and provide them with a learning environment—lectures, demonstrations, labs, and problems—that takes them from where they are to where we want them to be. Since we all know that a few

[32] See the papers referred to in the annotated bibliography [McDermott 1999] and in the review papers [McDermott 1991], [Reif 1994], and [Redish 1999].

[33] Audrey Champagne, private communication.

[34] We should not lose sight of the fact that the student's mental model in this case is in fact correct. We observe that lighter objects *do* fall more slowly than heavy ones if they fall in air, and few of us have much direct experience with objects falling in a vacuum. But for reasonably dense objects falling for only a few seconds, the difference is small, so that this observation does not yield a useful idealization. The observation that objects of very different mass fall in very nearly the same way does.

students get there from here using our current procedures, why can't we make it work for all of them? We do in fact know now that the right environment can produce substantially better physics learning in most of the students taking introductory university physics.[35] But my fourth principle is a word of warning that suggests we should not be looking for a "magic bullet."

> *Principle 4:* Since each individual constructs his or her own mental structures, different students have different mental responses and different approaches to learning. Any population of students will show a significant variation in a large number of cognitive variables.

I like to call this principle the *individuality* or *distribution function principle.* This reminds us that many variables in human behavior have a large natural line width. The large standard deviation obtained in many educational experiments is not experimental error; it is part of the measured result! As physicists, we should be accustomed to such data. We just aren't used to its being so broad and having so many variables. An "average" approach will miss everyone because no student is average in all ways.[36]

In addition to the fact that students have different experiences and have drawn different conclusions from them, their methods of approach may differ significantly. I state this as a corollary.

> *Corollary 4.1:* People have different styles of learning.

There is by now a vast literature on how people approach learning differently. Many variables have been identified on which distributions have been measured. These include authoritarian/independent, abstract/concrete, and algebraic/geometric, to name a few. The first variable means that some students want to be told, while others want to figure things out for themselves. The second means that some students like to work from the general to the specific, and some the other way round. The third means that some students prefer to manipulate algebraic expressions, while others prefer to see pictures. Many of us who have introduced the computer in physics teaching have noted that some students want to be guided step by step; others explore everything on their own. These are only a few of the variables. For some good analyses of individual cognitive styles and differences, see [Gardner 1999], [Kolb 1984], and [Entwistle 1981].

Once we begin to observe these differences in our students, we have to be exceedingly careful about how we use them. A preference does not mean a total lack of capability. Students who prefer examples with concrete numbers to abstract mathematical expressions may be responding to a lack of familiarity with algebra rather than a lack of innate ability. Many of our students' preferences come from years of being rewarded for some activities (such as

[35] As examples, see the references cited in the section "Research-Based Instructional Materials" in [McDermott 1999]. This annotated bibliography is also included on the Resource CD.

[36] This is analogous to the story of the three statisticians who went hunting deer with bows and arrows. They came across a large stag and the first statistician shot—and missed to the left. The second statistician shot and missed to the right. The third statistician jumped up and down shouting "We got him!"

being good memorizers) and chastised for others (such as asking questions the teacher couldn't answer). Expanding our students' horizons and teaching them how to think sometimes requires us to overcome years of negative training and what they themselves have come to believe are their own preferences and limitations.

An important implication is the following:

> *Corollary 4.2:* There is no unique answer to the question: What is the best way to teach a particular subject?

Different students will respond positively to different approaches. If we want to adopt the view that we want to teach all our students (or at least as many as possible), then we must use a mix of approaches and be prepared that some of them will not work for some students. We need to answer the question: What is the distribution function of learning characteristics that our students have in particular classes? Although some interesting studies have been done over the years, the implication for instruction in physics is not well understood.[37]

Another implication that is very difficult to keep in mind is:

> *Corollary 4.3:* Our own personal experiences may be a very poor guide for telling us the best way to teach our students.

Physics teachers are an atypical group. We "opted in" at an early stage in our careers because we liked physics for one reason or another. We then trained for up to a dozen years before we started teaching our own classes. This training stretches us even farther from the style of approach of the "typical" student. Is it any wonder that we don't understand most of our beginning students and they don't understand us?

I vividly recall a day a few years ago when a student in my algebra-based introductory physics class came in to ask about some motion problems. I said: "All right, let's get down to absolute basics. Let's draw a graph." The student's face fell, and I realized suddenly that a graph was not going to help him at all. I also realized that it was going to be hard for *me* to think without a graph and to understand what was going through the student's mind. I never minded doing without derivatives—motion after all is the study of experimental calculus, and you have to explain the concept (maybe without using the word "derivative") even in a non-calculus-based class. But I can't remember a time when I couldn't read a graph, and I have found it difficult to empathize with students who come to physics and can't read a graph or reason proportionately. It takes a special effort for me to figure out the right approach.

This is very natural given the earlier principles. Our own schemas for how to learn come from our personal reactions to our own experiences. However, to reach more of our students than the ones who resemble ourselves, we will have to do our best to get beyond this mindset. It makes the following principle essential.

> *Corollary 4.4:* The information about the state of our students' knowledge is contained within them. If we want to know what they know, we not only have to ask, we have to listen!

[37] See [Kolb 1984].

One point I want to stress about the individuality principle is in the idea expressed by its final sentence: *Any population of students will have a significant variation in a large number of cognitive variables.* We have a tendency, especially in physics, to classify students along a single axis that we think of as "intelligence." I've heard Sagredo say, "Well, most of my students are having trouble, but the smart ones get it." In cognitive science, there have been vigorous arguments for a number of years now as to whether there is a single variable (referred to as "g") that describes intelligence, or whether what we call intelligence consists of a number of independent factors.[38] The literature on this subject is quite complex, and I do not pretend to be able to evaluate it. However, whether or not intelligence is a unary concept, success in physics—or in any scientific career—relies on much more than intelligence. I have followed with interest the careers of many of my physics classmates from college and graduate school for many decades, and one thing is absolutely clear. The students who were the "brightest" in doing their schoolwork were not necessarily the ones who went on to make the most important contributions to physics. Creativity, persistence, interpersonal skills, and many other factors also played large roles.[39] This point is discussed again later in the next chapter on the hidden curriculum.

5. The social learning principle

With the fifth principle, I go beyond single individuals and consider their relations to others as a part of their learning. This principle is based on the work on group learning that builds on the ideas of the Russian psychologist, Lev Vygotsky. These ideas have had a profound impact on modern theories of teaching and learning [Vygotsky 1978] [Johnson 1993].

> *Principle 5:* For most individuals, learning is most effectively carried out via social interactions.

The social learning principle is particularly important for physicists to keep in mind. Physicists as a group are highly unusual in many ways. In my experience, they tend to be in the extreme tails of distributions of curiosity, independence, and mathematical skills. They also tend to be highly self-sufficient learners. I once heard David Halliday remark that what he enjoyed most as a student was sitting down by himself alone in a quiet room with a physics text and going "one-on-one" with the authors of the book—trying to understand them and figure out what they were saying. Many of us have similar inclinations. Physicists as a group seem to be selected for the character of being able to learn on their own. But in examining my experiences of this type, I have decided that my "learning on my own" involves an ability to create an "internalized other"—to take a variety of viewpoints and to argue with myself. This is not a commonly found characteristic and should not be assumed in a general population of students.

[38] For a discussion and for references to this literature, see [Gardner 1999].

[39] The range of important variables provides a basis for what we might call *Garrison Keillor's Corollary.* "All students are above average—on some measure."

SOME GENERAL INSTRUCTIONAL METHODS DERIVED FROM THE COGNITIVE MODEL

Many instructional methods have been developed based on the cognitive models discussed above. Two that will be relevant for us later in our discussions of specific curricula are *cognitive conflict* and *bridging*. The cognitive conflict method is used when an inappropriate generalization or incorrect association has become particularly robust and difficult to change. The bridging method relies on the explicit idea that the students bring useful knowledge resources to their learning of physics and attempts to explicitly activate those resources in appropriate ways.

Cognitive Conflict

Common naïve conceptions can be strikingly robust. At the beginning of my unit on direct current circuits in my calculus-based engineering physics class, I gave the students the problem shown in Figure 2.7.

This problem is by now rather traditional. It appears in sixth grade science texts and in many high school physics texts. My discussion of circuits occurred at the beginning of our third semester of engineering physics. Most of the students were second-semester sophomores, and many of them were electrical engineering students. More than 95% of them had taken (and done well in) high school physics. Yet only about 10-15% of them were able to solve this problem before instruction.

Sagredo complained that this was "a trick question." He said, "I'll bet many graduate students in physics would miss it. You have to find a clever way to make it light—touching the bulb's tip to one end of the battery to make a contact without a wire. Then you can use the wire to close the circuit."

Sagredo's correct answer is shown at the left of Figure 2.8. He is right that many physics grad students will miss the problem on the first try (and many professors as well), but the specific wrong answers given by students and by experts show that something different is going on. The experts who get it wrong give the answer, "No. It can't be done. You need a closed circuit and that takes two wires." My engineering students' wrong answers were, "Sure you

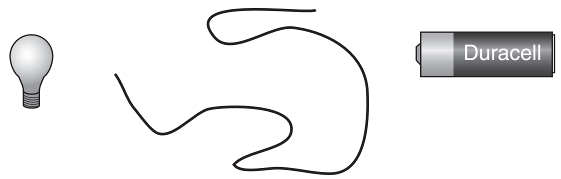

Given the three objects shown above, a battery, a bulb, and a single piece of wire, can you make the bulb light without cutting the wire? If you can, show how. If you can't, explain why not.

Figure 2.7 A problem introductory students often have difficulty with.

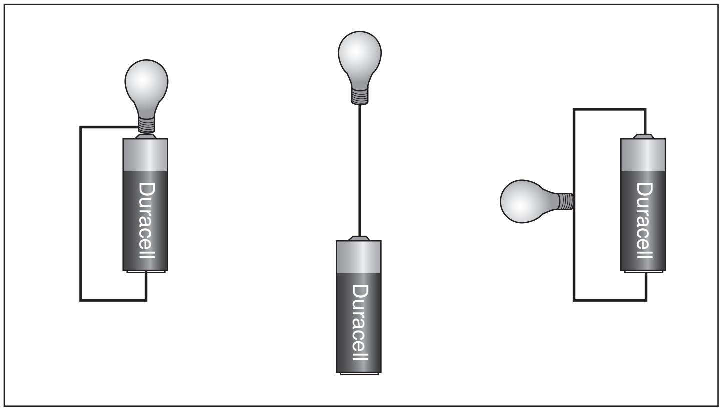

Figure 2.8 Correct answer (on left) to battery-bulb-wire problem and the two most common incorrect student answers.

can. It's easy." One-third of the students gave the answer shown in the middle, and one-third gave the answer shown at the right. About half of the rest gave the correct answer. The rest either left it blank or gave uninterpretable answers. Very few students gave the "expert's error."

Many students showed the common naïve conception of electricity as a source of energy that can be "drawn" directly from an electric power source or "tapped into" by touching. Students' naïve conceptions about electric currents are documented in the research of the University of Washington Physics Education Group (UWPEG) [McDermott 1992]. The group also reports a lesson they have developed to address this issue [Shaffer 1992]. The lesson is delivered in a "Tutorial"—a structure that has replaced recitations at the University of Washington. These and related materials are a part of the Physics Suite and are discussed in detail in chapter 8.

The model frequently used in Tutorials is cognitive conflict in the form *elicit/confront/resolve/reflect*. In the first Tutorial on direct currents, the lesson begins with the question shown above (given on a pretutorial ungraded quiz[40] during lecture). When the students get to the Tutorial, each group gets a battery, a bulb, and a single wire.

When I gave this lesson in my classes, about half of the students expected to be able to light the bulb using one of the two arrangements at the right of Figure 2.8. I particularly remember one student who came up to me complaining that the equipment she had been given had to be defective. (She was certain that the middle arrangement shown in Figure 2.8 should work. "After all," she said, "that's the way a flashlight works, isn't it?") She insisted on having a new bulb and a fresh battery.

The subsequent discussions with their colleagues and with the facilitators (and the rest of the lesson, which elaborates on the point and reconsiders it in a variety of contexts) help students to resolve the conflict between their model and their observations, and to reflect

[40]One reason for not grading the quiz is to encourage students to look for what they actually think is plausible rather than to try to guess what the teacher wants. The solutions also are not posted since the point is to get students thinking about the reasoning (during the Tutorial period) rather than focusing on the answer.

upon the implications. After Tutorial instruction, students did significantly better (~75% success) on other (more complex) questions probing similar issues than students with traditional instruction (< 50% success). Many other examples exist of successful learning produced through lessons based on cognitive conflict.

Bridging

The use of cognitive conflict as the primary instructional tool in the classroom can lead to difficulties because it is a rather negative approach. A colleague who had instituted a reform procedure in a high school physics class that relied heavily on the cognitive-conflict method reported to me that one of his students responded to a quiz: "Oh great. Here's another test to show us how stupid we are about physics." Once students are convinced they cannot do physics, it is extremely difficult to teach them anything more than rote memorization.[41]

However, if we consider the student's starting knowledge as a resource and begin with what is *right* about student reasoning, we can make them feel considerably better about themselves and about their ability to do physics. To see how this might work, let's consider an example.

John Clement has proposed building on the correct aspects of students' naïve conceptions by finding a series of *bridges* or *interpolating steps* that help students transform their mental models to match the accepted scientific one [Clement 1989]. The starting point for a bridge should be something that the students know well that is correct—which Clement refers to as an *anchor*. An example of bridging is in Clement's building of a lesson to respond to the common naïve conception about normal forces.

In his classic paper, "Explaining the 'At Rest' Condition of an Object," Jim Minstrell documented that more than half of the students in his high school physics class did not believe that an inanimate supporting object could exert a force [Minstrell 1982]. They interpreted what a table does to a book resting upon it as "simply preventing it from falling." They showed that the pattern of association triggered by the word "force" had a strong link to the idea of will or intent and hence, to something produced by an "active" object, most frequently an animate one. Students seem to be bringing up a "blocking" reasoning primitive rather than a "springiness" one. This result has subsequently been confirmed to occur in a significant fraction of introductory physics students.

Clement looked for examples in which students had correct intuitions (primitive responses or "gut feelings") that could be used to build on. He suggests that a useful starting point for a bridging lesson should have the following characteristics.

1. It should trigger with high probability a response that is in rough agreement with the results of correct physical theory.
2. It should be concrete rather than abstract.
3. The student should be internally confident of the result; that is, students should strongly believe in the reasoning by themselves and not by virtue of authority.

Clement came up with two possible anchors for the situation described by Minstrell.

- Holding up a heavy book in your hand
- An object being held up by a reasonably soft spring

[41] These issues are discussed in more detail in chapter 3.

TABLE 2.1 Results of a Bridging Lesson

	Pre-test average	Post-test average	Fraction of possible gain achieved
Control group ($N = 55$)	(17 ± 1)%	(45 ± 2)%	0.34
Experimental group ($N = 155$)	(25 ± 2)%	(79 ± 1)%	0.72

To Clement's surprise, the spring served as a much better starting point than the book. This stresses the fact that we cannot assume that an example that seems an obvious anchor to us will work for our particular population. They must be tested. Our individuality principle also reminds us not to expect that a particular example will serve for all students. We may have to use a number of different anchors to reach most students.

Clement tested this approach in four high schools in Massachusetts [Clement 1998]. There were 155 students in the experimental group and 55 in the control. The same pre- and post-tests were given to all students. On the issue of static normal forces, the gains of the groups were as shown in Table 2.1. The experimental group was taught with a bridging lesson; the control class with a traditional one. Evaluation was done with six questions that were part of a longer test. (Errors are std. error of the mean.) Similar results were reported for bridging strategies in the areas of friction and dynamic use of Newton's third law in collisions.[42]

Restricting the frame

The real world, with all its complexity, is too much for us to handle all at once. Long-term and working memories function together to parse the visible world into the pieces relevant for functioning in any particular situation. Many students, when facing physics problems, have considerable difficulty figuring out what is important and what is not. Helping them build appropriate templates and associations is an important part of what physics instruction is trying to accomplish.

Limiting our view of what we want to analyze is an essential part of the scientific enterprise. One of Galileo's greatest contributions to science was his ability to step back from the Aristotelian attempt to produce a grand synthesis for all motion and to focus on understanding well how a few simple phenomena really worked—an object on an inclined plane, the pendulum, a falling body. The synthesis comes later, once you have some solid bricks to build into a more coherent structure. When Newton synthesized a theory of motion 50 years later, he created a limited synthesis—a theory of motion, but not a single theory that successfully described, for example, light, heat, and the properties of matter at the same time. The scientific structure grows at different paces and in different places, with pieces being continually matched and modified to create more coherent and useful maps.

Similarly, in teaching introductory physics, we also have to restrict our considerations to a piece of the entire picture. Our goal is not to present the most coherent picture of the entire physical knowledge system that we as professionals have been able to construct, even

[42]For a discussion of the lessons and the controls, see [Clement 1998].

though this might be an intellectually enticing and enjoyable goal. The mental structure that we've created for ourselves to describe physics has been built up over many years—perhaps decades—and has resulted from continual transformations that have been made not only to the physics, but to ourselves.

Even in solving a single problem in introductory physics, we must go through the process of restricting consideration to a piece of the world and of limiting those aspects we want to consider. I refer to selecting a limited piece of the world to view through a frame as *cat television*. My cat (see Figure 2.9) very much enjoys viewing a small piece of the world through a window (as have many other cats I have known) and can become quite addicted to it. When we drove across the country for a sabbatical visit with the cat in the car, in every motel room he insisted on finding a place where he could look out the window and see where in the world we were. The real TV doesn't interest him at all. (The same is true for many physicists.)

Once we have drawn our frame—chosen our channel on cat TV—we still have work to do before we can start drawing a representation. We have to choose what to pay attention to and what to ignore. I call this process *creating the cartoon*. I mean "cartoon" here in the sense of an artist's sketch, drawn in preparation to creating a painting or mural. The process of deciding what to keep and what to ignore is a difficult one, and one that is often glossed over by physicists who already know what to consider. (An example is Bill Amend's FoxTrot cartoon—in the other sense—shown in Figure 2.10.)

Changing what one decides to look at in a real-world physical phenomenon can be associated with a major advance in science. In his *Structure of Scientific Revolutions*, T. S. Kuhn [KuhnT 1970] describes the paradigm shift associated with Galileo's observation of the pendulum. The story goes that Galileo was sitting in church, bored with an extended sermon, perhaps, when a gust of wind started a chandelier swinging. An observation that could be made on the chandelier is that it eventually comes to rest. One could infer that this supports the Aristotelian position that objects naturally come to rest (and seek the lowest point). But Galileo timed the swing of the chandelier using his pulse and noted that the period of the

Figure 2.9 "Cat television." Cats often enjoy observing a small, framed "piece of the world." Similarly, a scientist must narrow his or her focus in a particular set of issues in order to make progress. I call this "choosing a channel on cat TV."

Figure 2.10 Assumptions in traditional physics problems are often tacit but important.

swing did not change as the amplitude got smaller. He then realized that this could be interpreted as saying that the natural state of the object was to continue to oscillate. In the Aristotelian view, the coming to rest is taken as fundamental, and the oscillation is seen as a deviation. In the Galilean view, the oscillation is seen as fundamental, and the damping is seen as peripheral. There is no *a priori* way to decide which view is better. One must have more information as to how these assumptions play out and how productive they are in generating deeper understandings of a wider class of phenomena.

When students view a subject or a problem in physics, they are often unclear as to how to link what the physicist sees in the problem with their personal experience. A real-world problem involves friction, massive pulleys, cartwheels that have significant moments of inertia, etc., etc., etc. *Knowing what to ignore is an important part of learning to think about science, and it should not be treated as trivial.* Sagredo often teaches his students "the first step in doing a physics problem is to draw a diagram," but expresses frustration that often the students "can't even do that." I have seen students who, when given that instruction, draw a 3-D diagram when a 2-D one will do, or who carefully draw a realistic car, person, or horse when the problem could be done by using a block. Is it obvious, when asked to "consider a car rolling down a hill," that it doesn't matter whether the car is a Porsche or a VW bug or a block of wood with wheels? I conjecture that our ignoring this step is one reason that students don't always see the relation of physics to their real-world experience (see chapter 3).

Multiple representations

In science we use a dizzying variety of representations[43] that may have a variety of positive cognitive effects. First, the use of multiple representations (visual and verbal, for example) may help us make better use of our working memory. Second, different representations associate more naturally with different features of the data/situation we are trying to describe. As a result, the use of multiple representations can be effective in building links between different aspects of the situation. Some of the representations we use in physics include

- words
- equations
- tables of numbers
- graphs
- specialized diagrams

For the expert, each representation expresses some characteristic of the real-world system more effectively than others, and the translation and links between them help build up a strong and coherent mental model. But each representation takes some time and effort to learn. If students haven't yet learned to "read" a representation, that is, if their knowledge of that representation hasn't been compiled into an easily processed chunk, the translation of the representation can take up too much working memory to allow the students to make effective use of the representation. Some may have mastered one or more of these representations through experiences they had before coming to our class. Some may have a learning style that favors a verbal or a visual representation. Some may think they have trouble with one or more of our representations and actively avoid thinking about them [Lemke 1990].

I represent our process of setting up a problem or scientific exploration with the diagram shown in Figure 2.11. We begin by picking a channel of cat television with a specific real-world situation considered in a specific physics context—for example, the mechanics of a car rolling down a hill. We then make a selection of what we want to look at first—a simplified image that corresponds to a box sliding down a frictionless inclined plane—what I called above: creating a cartoon. Is this a good model for an SUV rolling down a rocky mountain trail? Perhaps not. But it might be a good place to start for a car with smaller wheels rolling down a paved road.

Once we have our cartoon, we then express our knowledge of the situation in a variety of ways, using linked multiple representations. These present the information in different ways, enabling different ways of thinking about and linking the information.

There is considerable evidence in the research literature documenting student difficulties with representation translation [Thornton 1990] [Beichner 1994]. Students often see, for example, the drawing of a graph as the solution to a problem—something the teacher asked them to do—instead of a tool to help them understand and solve a more complex problem.

[43] See, for example, the detailed analysis of a chemistry and physics class in [Lemke 2000].

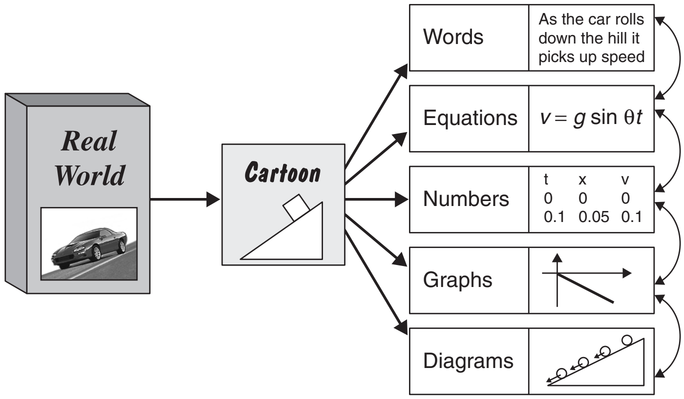

Figure 2.11 Physicists use a large number of different representations to describe events and processes in the real world. Learning to handle these representations and to translate them is an important—and difficult—step for introductory physics students.

In order to use multiple representations in our physics classes, it is important that we be aware of their difficulty for students and that we help those students both to learn the physics and to learn to make the connection to the modes of thinking they are less comfortable with.

RETHINKING THE GOALS OF PHYSICS INSTRUCTION

Putting our students' learning into a cognitive framework helps us begin a more detailed and constructive dialog in the community of physics teachers on what we want to accomplish with our physics instruction. There is more to learning physics than placing check marks on the table of contents of a text. Many of the most important results of our instruction are not associated with particular physics content. Even those goals that are associated with particular content can now be seen in a different way, given our understanding of student thinking. In this section, I try to explicate some of the "hidden" goals of attitudes and skills that we might like our students to attain in an introductory physics course—goals that are rarely discussed and whose attainment is often both strongly desired and taken for granted.

"Wait!" Sagredo interjects. "This talk of attitudes and skills is all very well, and I hope my students will develop them. But the physics content is important and we shouldn't lose sight of it." An excellent point, Sagredo. Let's start with what our learning theory says about learning physics content itself. Then we can consider attitude goals and skill development.

Extended content goals

Learning physics content means much more than memorizing a lot of independent definitions and equations. We want our students to understand enough about what the physics means to be able to understand what problems are about and what their answers imply; we want them to understand how the physics they are learning fits—both with other physics they have learned and with their personal experience with the world; and we want them to be able to use their knowledge effectively in solving problems. I refer to these three goals as *concepts*, *coherence*, and *functionality*.

> *Goal 1:* Concepts—Our students should understand what the physics they are learning is about in terms of a strong base in concepts firmly rooted in the physical world.

From our cognitive model we know that students can attach new knowledge to their existing knowledge structures as something separate and only relevant for the context of solving problems in a physics class. We want students to not only compile the mathematical manipulations associated with solving physics problems (i.e., learn to use them without thinking about them); we want them to understand what the physics is *about*.

To achieve this goal, students have to make sense of physics concepts—the ideas and definitions that map abstract physics descriptions onto real things in the physical world. In order to help them reach this goal through instruction, it often helps to motivate the need for a concept before introducing the definition through direct observation of phenomena. Arnold Arons called this "idea first, name afterwards" and stressed the value of operational definitions in physics instruction. He has a nice description of this in his book [Arons 1990].

> In my own courses, I indicate from the first day that we will operate under the precept "idea first and name afterwards," and that scientific terms acquire meaning only through the description of shared experience in words of prior definition. When students try to exhibit erudition (or take refuge from questioning) by name dropping technical terms that have not yet been defined, I and my staff go completely blank and uncomprehending. Students catch on to this game quite quickly. They cease name dropping and begin to recognize, on their own, when they do not understand the meaning of a term. Then they start drifting in to tell me of instances in which they got into trouble in a psychology or sociology, or economic, or political science course by asking for the operational meaning of technical terms. [Arons 1990]

While presenting "idea first, name second" is a good start, it's rarely sufficient to get students to develop a good understanding of physics concepts. A considerable amount of the effort in physics education research over the past decade has been devoted to the development of effective instructional techniques for helping students build their conceptual understanding. For more details on these specific content issues, see [Arons 1990] and the references in the Resource Letter [McDermott 1999] given on the resource CD.

> *Goal 2:* Coherence—Our students should link the knowledge they acquire in their physics class into coherent physical models.

A major strength of the scientific worldview is its ability to describe many complex phenomena with a few simple laws and principles. Students who emphasize science as a collection of facts may fail to see the integrity of the structure, an integrity that is both convincing and useful in solving problems. The lack of a coherent view can cause students many difficulties, including a failure to notice errors in their reasoning and an inability to evaluate a recalled item through cross-checks.

Let's recall Principles 1 and 2: Students will put what we give them into their knowledge structure and integrate it into their existing knowledge structure in some way of their own. Whatever it is that they know of any particular content, they may or may not create coherent schemas with appropriate connections that are activated with a high probability in appropriate contexts. We want our students to not simply "get the content" but to build their understanding of that content into an accurate and effective mental model.

> *Goal 3:* Functionality—Our students should learn both *how* to use the physics they are learning and *when* to use it.

For most of our populations of introductory physics students, we don't just want them to *know* the physics content but to be able to *do something with it.* My second cognitive principle suggests that in addition to having students master the physics content and see that it makes sense, we also want students' knowledge of physics to be robust and functional. That is, they should be able to recognize contexts that are appropriate for their particular physics knowledge and use them in correct ways. This means that we need to help students not only to obtain physics content knowledge, but to organize their knowledge.

These goals suggest that we should broaden our evaluation procedures. Traditionally we only test the content and part of the student's skill in doing physics, usually (at least at the introductory level) in limited preset contexts. Sagredo once decided that an effective way to get students to do their homework was to create exams made up of only previously assigned homework problems. The students were quite happy with this arrangement. Unfortunately, this sends the strong message that the only physics knowledge they need is restricted to a small set of problems whose solutions can be memorized. Students trained in this fashion are unlikely to be able to solve any other problems than the ones they have memorized. (I have seen situations where reversing the figure from left to right made the problem impossible for students who were quite comfortable solving the original problem.)

It is not sufficient for students to "know" the relevant correct statements of physics. They also have to be able to gain access to them at the appropriate times; and they have to have methods of cross-checking and evaluating to be certain that the result they have called up is truly relevant. To do this, they need to build a coherent and consistent mental model of the subject.

If we want to help our students build good models of the physics content, our cognitive model of student learning provides some guidance. The experience that outstanding teachers have reported is consistent with what we learn from researchers in cognitive and neuroscience: activities teaching thinking and reasoning have to be repeated in different contexts over a period of time. Arnold Arons observes:

> It must be emphasized, however, that *repetition* is an absolute essential feature of [effective] instruction—repetition *not* with the same exercises or in the same context but in continually altered and enriched context. . . . Experience . . . must be spread out over weeks and months and must be returned to in new contexts after the germination made possible by elapsed time. Starting with only a few students being successful, each repetition or re-cycling sees an additional percentage of the class achieving success, usually with a leveling of somewhat below 100% of the total after approximately five cycles. [Arons 1983]

Our cognitive colleagues tell us that to move something from short-term to long-term memory takes repetition and rehearsal. Our neuroscience colleagues are showing us that learning involves real structural changes among neural connections. I summarize this with a basic teaching and learning precept—really a corollary to the constructivism principle.

> *Corollary 1.1:* Learning is a growth, not a transfer. It takes repetition, reflection, and integration to build robust, functional knowledge.

This corollary leads to a guideline for instruction. I refer to it (and to subsequent guidelines for instruction) as a "teaching commandment." The full set (along with the cognitive principles and goals listed in this chapter and elsewhere) are summarized in a file on the Resource CD.

> *Redish's first teaching commandment:* Building functional scientific mental models does not occur spontaneously for most students. They have to carry out repeated and varied activities that help build coherence.