DANIEL T. WILLINGHAM

WHY DON'T STUDENTS *Like* SCHOOL?

A COGNITIVE SCIENTIST ANSWERS QUESTIONS ABOUT HOW THE MIND WORKS AND WHAT IT MEANS FOR THE CLASSROOM

3

Why Do Students Remember Everything That's on Television and Forget Everything I Say?

Question: Memory is mysterious. You may lose a memory created fifteen seconds earlier, such as when you find yourself standing in your kitchen trying to remember what you came there to fetch. Other seemingly trivial memories (for example, advertisements) may last a lifetime. What makes something stick in memory, and what is likely to slip away?

Answer: We can't store everything we experience in memory. Too much happens. So what should the memory system tuck away? Things that are repeated again and again? But what about a really important one-time event such as a wedding? Things that cause emotion? But then you wouldn't remember important yet neutral things (for example, most schoolwork). How can the memory system know what you'll need to remember later? Your memory system lays its bets this way: if you think about something carefully, you'll probably have to think about it again, so it should be stored. Thus your memory is not a product of what you want to remember or what you try to remember; it's a product of what you think about. A teacher once told me that for a fourth-grade unit on the Underground Railroad he had his students bake biscuits, because this was a staple food for runaway slaves. He asked what I thought about the assignment. I pointed out that his students probably thought for forty seconds about the relationship of biscuits to the Underground Railroad, and for forty minutes about measuring flour, mixing shortening, and so on. Whatever students think about is what they will remember. The cognitive principle that guides this chapter is:

> Memory is the residue of thought.

To teach well, you should pay careful attention to what an assignment will actually make students think about (not what you hope they will think about), because that is what they will remember.

The Importance of Memory

Every teacher has had the following experience: you teach what you think is a terrific lesson, full of lively examples, deep content, engaging problems to solve, and a clear message, but the next day students remember nothing of it except a joke you told and an off-the-subject aside about your family[1]—or worse, when you say, struggling to keep your voice calm, "The point of yesterday's lesson was that one plus one equals two," they look at you incredulously and say, "One plus one equals *two?*" Obviously, if the message of Chapter Two is "background knowledge matters," then we must closely consider how we can make sure that students acquire this background knowledge. So why do students remember some things and forget other things?

Let's start by considering why you fail to remember something. Suppose I said to you, "Can you summarize the last professional development seminar you attended?" Let's further suppose that you brightly answer, "Nope, I sure can't." Why don't you remember?

One of four things has happened, all of which are illustrated in Figure 1, a slightly elaborated version of the diagram of the mind that we've used before. You will recall that working memory is where you keep things "in mind," the location of consciousness. There is lots of information in the environment, most of which we are not aware of. For example, as I write this, the refrigerator is humming, birds are chirping outside, and there is pressure on my backside from the chair I'm sitting on—but none of

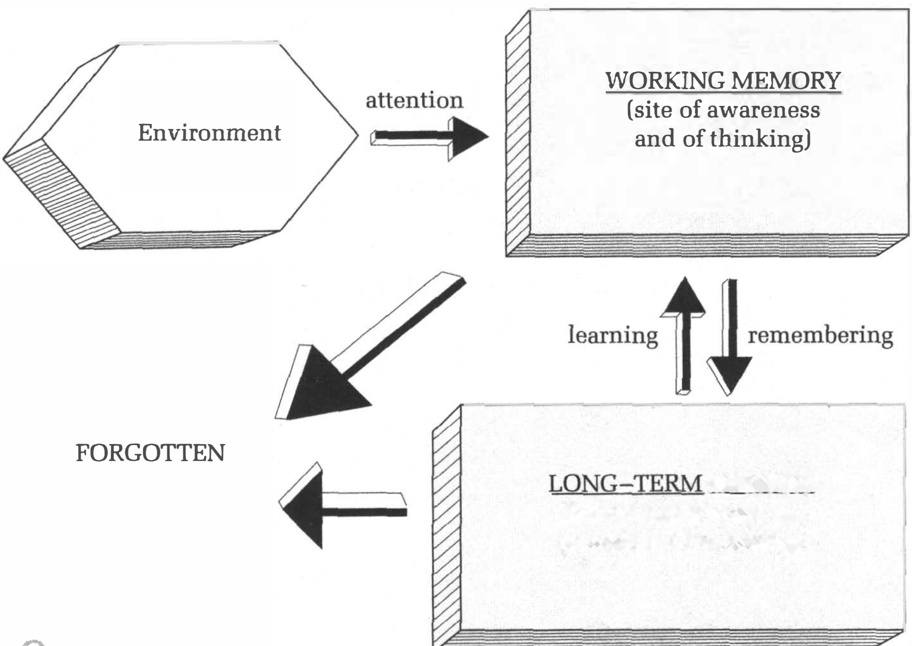

FIGURE 1: A slightly modified version of our simple diagram of the mind.

that was in my working memory (that is, my awareness) until I paid attention to it. As you can see in Figure 1, things can't get into long-term memory unless they have first been in working memory. So this is a somewhat complex way of explaining the familiar phenomenon: *If you don't pay attention to something, you can't learn it!* You won't remember much of the seminar if you were thinking about something else.

Information can enter working memory not only from the environment but also from long-term memory; that's what I mean when I refer to remembering, as shown by the labeled arrow. So another possible reason you don't remember is that the process by which things are drawn from long-term memory has failed. I discuss why that happens in Chapter Four.

A third possibility is that the information no longer resides in long-term memory—that it has been forgotten. I'm not going to discuss forgetting, but it's worth taking a moment to dispel a common myth. You sometimes hear that the mind records in exquisite detail everything that happens to you, like a video camera, but you just can't get at most of it—that is, memory failures are a problem of access. If you were given the right cue, the theory goes, anything that ever happened to you would be recoverable. For example, you may think you remember almost nothing of your childhood home, but when you revisit it the smell of the camellia blooms in the yard wipes away the years, and the memories that you thought were lost can be pulled out, like charms on a fine chain. Such experiences raise the possibility that *any* memory that you believe is lost can in principle be recovered again. Successful memory under hypnosis is often raised as evidence to support this theory. If the right cue (camellia blossoms or whatever it might be) can't be found, hypnosis allows you to probe the vault directly.

Although this idea is appealing, it's wrong. We know that hypnosis doesn't aid memory. That's easy to test in the laboratory. Simply give people some stuff to remember, then later hypnotize half of them and compare their recall to that of the people who were not hypnotized. This sort of experiment has been done dozens of times, and typical results are shown in Figure 2.[2] Hypnosis doesn't help. It does make you more confident that your memory is right, but it doesn't actually make your memory more accurate.

The other bit of evidence—that a good cue such as the odor of camellia can bring back long-lost memories—is much more difficult to test in a laboratory experiment, although most memory researchers believe that such recoveries are possible. But even if we allow that lost memories can be recovered in this way, it doesn't mean that *all* seemingly forgotten memories are recoverable—it just means that a few are. In sum, memory researchers see no reason to believe that all memories are recorded forever.

Now, let's return to our discussion of forgetting. Sometimes you *do* pay attention, so the material rattles around working memory for a while, but it never makes it to long-term memory. An example of a few such bits of information from my own experience are shown in Figure 3. *Lateral line* is a term I have looked up more than once, but I couldn't tell you now what it means. You doubtless have your own examples of things you are certain you *ought* to know, because you've looked them up or heard them (and thus they have been in working memory), yet they have never stuck in long-term memory.

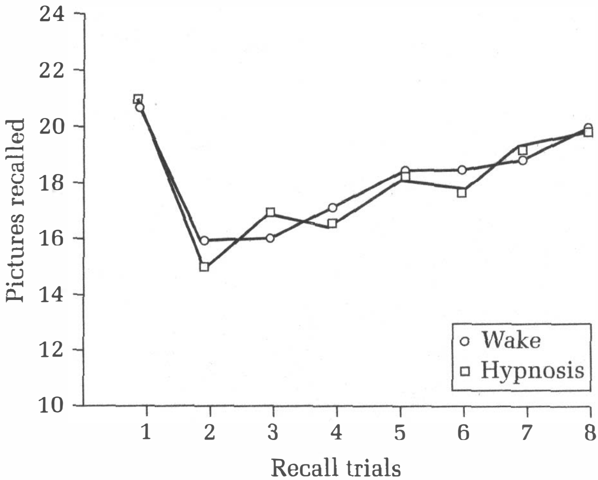

FIGURE 2: Subjects were shown forty drawings of common objects and then had to try to recall them. Session 1 happened right away; sessions 2 through 8 occurred a week later. Naturally there was significant forgetting during the week, and with each attempt to remember, subjects on average did recall more. Also, the hypnotized subjects didn't remember any more than the nonhypnotized subjects.

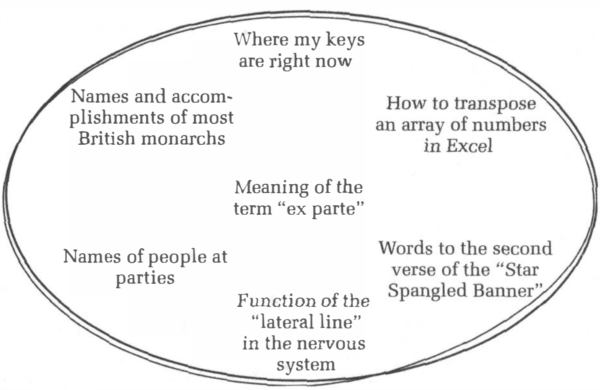

FIGURE 3: Bits of information that I am certain I have paid attention to and that thus have resided in my working memory but that have never made it into my long-term memory.

Just as odd is that some things have remained in your long-term memory for years although you had no intention of learning them; indeed, they held no special interest for you. For example, why do I know the jingle from the 1970s Bumble Bee tuna advertisement (Figure 4)?

You could make a good argument that understanding the difference between Figure 3 and Figure 4 is one of the core problems in education. We all know that students won't learn if they aren't paying attention. What's more mysterious is why, when they *are* paying attention, they sometimes learn and sometimes don't. What else is needed besides attention?

A reasonable guess is that we remember things that bring about some emotional reaction. Aren't you likely to remember really happy moments, such as a wedding, or really sad ones, such as hearing the news of the attacks on 9/11? You are, and in fact if you ask people to name their most vivid memories, they often relate events that probably had some emotional content, such as a first date or a birthday celebration (Figure 5).

Naturally we pay more attention to emotional events, and we are likely to talk about them later, so scientists have had to conduct very careful studies to show that it's really the emotion and not the repeated thought about these events that provides the boost to memory. The effect of emotion on memory is indeed real, and researchers have actually worked out some of the biochemistry behind it, but the emotion needs to be reasonably strong to have much impact on memory. If memory *depended* on emotion, we would remember little of what we encounter in school. So the answer *Things go into long-term memory if they create an emotional reaction* is not quite right. It's more accurate to say, *Things that create an emotional reaction will be better remembered, but emotion is not necessary for learning.*

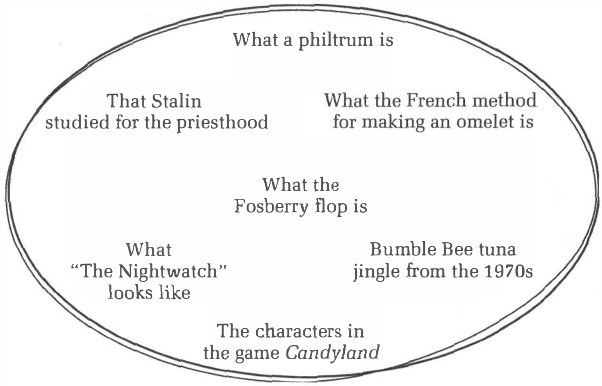

FIGURE 4: Material that is in the author's long-term memory even though the author didn't want to learn it and was in fact not all that interested in it.

Repetition is another obvious candidate for what makes learning work. Maybe the reason I remember the Bumble Bee tuna jingle (Figure 4) from thirty years ago is that I heard it a lot. Repetition is very important, and I discuss it in Chapter Five, but it turns out that not just any repetition will do. Material may be repeated

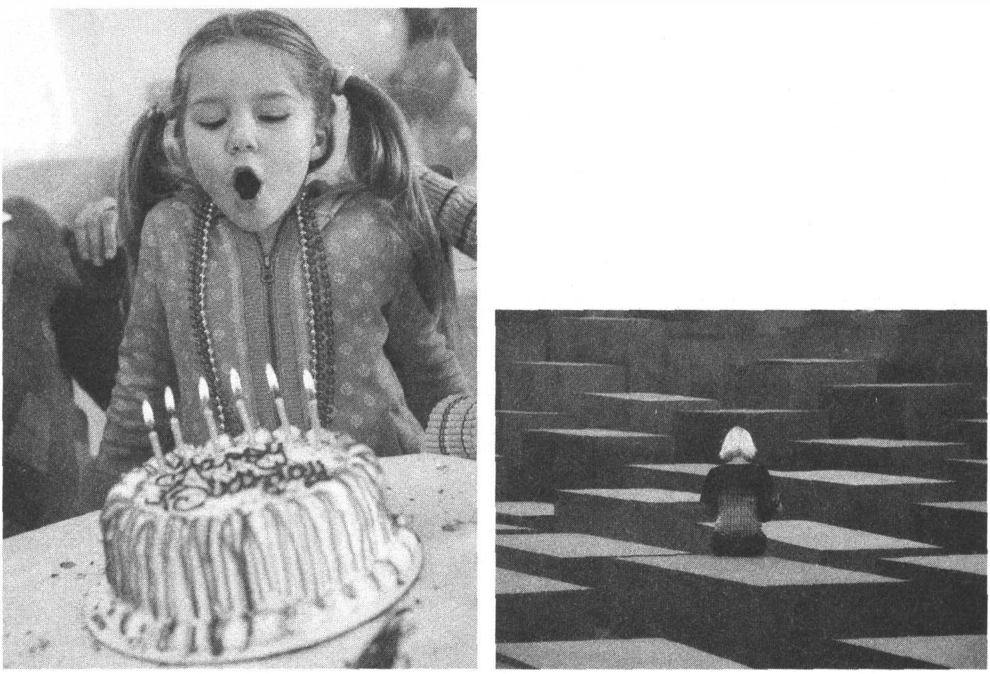

FIGURE 5: Emotional events tend to be well remembered, whether they are happy, such as a birthday party, or sad, such as a visit to the Holocaust Memorial in Berlin.

almost indefinitely and still not stick in your memory. For example, have a look at Figure 6. Can you spot the real penny among the counterfeits?

You have seen thousands of pennies in your lifetime—a huge number of repetitions. Yet, if you're like most people, you don't know much about what a penny looks like.[3] (The real penny is choice *A,* by the way.)

So repetition alone won't do it. It's equally clear that *wanting* to remember something is not the magic ingredient. How marvelous it would be if memory did work that way. Students would sit down with a book, say to themselves, "I want to remember this," and they would! You'd remember the names of people you've met, and you'd always know where your car keys are. Sadly, memory doesn't work that way, as demonstrated in a classic laboratory experiment.[4] Subjects were shown words on a screen one at a time and were asked to make a simple judgment about each word. (Some subjects had to say whether the word contained either an *A* or a *B*; others had to say whether the word made them think of pleasant things or unpleasant things.) An important part of the experiment was that half of the subjects were told that their memory for the words would be tested later, after they had seen the whole list. The other subjects were not warned about the test. One of the remarkable findings was that knowing about the future test didn't improve subjects' memories. Other experiments have shown that telling subjects they'll be paid for each remembered word doesn't help much. So *wanting* to remember has little or no effect.

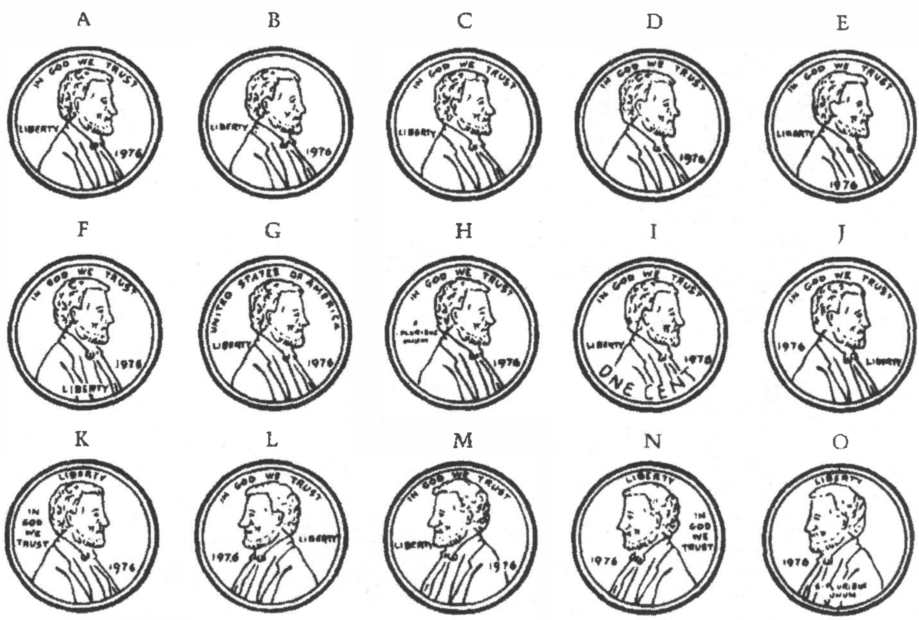

FIGURE 6: Can you find the real penny among the counterfeits? People are terrible at this task even though they have seen a penny thousands of times.

But there's another finding from this experiment that's still more important. Remember that when subjects saw each word they had to make a judgment about it—either about whether it contained an *A* or a *Q,* or about whether it made them think of pleasant or unpleasant things. The people who made the second type of judgment remembered nearly twice as many words as the people who made the first judgment. Now we seem to be getting somewhere. We've found a situation in which memory gets a big boost. But why would it help to think about whether a word is pleasant or not?

In this case it matters because judging pleasantness makes you think about what the word *means* and about other words that are related to that meaning. Thus, if you saw the word *oven,* you might think about cakes and roasts and about your kitchen oven, which doesn't work well, and so on. But if you were asked to judge whether *oven* contained an *A* or a *Q,* you wouldn't have to think about the meaning at all.

So it seems we're poised to say that *thinking about meaning is good for memory.* That's close, but not quite right. The penny example doesn't fit that generalization. In fact, the penny example shows just the opposite. I said that you've been exposed to a penny thousands of times (at least), and most of those times you were thinking about the penny's meaning—that is, you were thinking about its function, about the fact that it has monetary value, even if that value is modest. But having thought about the meaning of a penny doesn't help when you're trying to remember what the penny looks like, which is what the test in Figure 6 requires.

Here's another way to think about it. Suppose you are walking the halls of your school and you see a student muttering to himself in front of his open locker. You can't hear what he's saying, but you can tell from his tone that he's angry. There are several things you could focus on. You could think about the *sound* of the student's voice, you could focus on how he *looks,* or you could think about the *meaning* of the incident (why the student might be angry, whether you should speak to him, and so on). These thoughts will lead to different memories of the event the next day. If you thought only about the sound of the student's voice, the next day you'd probably remember that sound quite well but not his appearance. If you focused on visual details, then that's what you'd remember the next day, not what the student's voice sounded like. In the same way, if you think about the meaning of a penny but never about the visual details, you won't remember the visual details, even if they have been in front of your eyes ten thousand times.

Whatever you think about, that's what you remember. *Memory is the residue of thought.* Once stated, this conclusion seems impossibly obvious. Indeed, it's a very sensible way to set up a memory system. Given that you can't store everything away, how should you pick what to store and what to drop? Your brain lays its bets this way: If you don't think about something very much, then you probably won't want to think about it again, so it need not be stored. If you do think about something, then it's likely that you'll want to think about it *in the same way* in the future. If I think about what the student looks like when I see him, then his appearance is probably what I'll want to know about when I think about that student later.

There are a couple of subtleties to this obvious conclusion that we need to draw out. First, when we're talking about school, we usually want students to remember what things mean. Sometimes what things look like is important—for example, the beautiful facade of the Parthenon, or the shape of Benin—but much more often we want students to think about meaning. Ninety-five percent of what students learn in school concerns meaning, not what things look like or what they sound like.* Therefore, a teacher's goal should almost always be to get students to think about meaning.

The second subtlety (again, obvious once it's made explicit) is that there can be different aspects of meaning for the same material. For example, the word *piano* has lots of meaning-based characteristics (Figure 7). You could think about the fact that it makes music, or about the fact that it's expensive, or that it's really heavy, or that it's made from fine-quality wood, and so on. In one of my all-time favorite experiments, the researchers led subjects to think of one or another characteristic of words by placing them in sentences—for example, "The moving men lugged the PIANO up the flight of stairs" or "The professional played the PIANO with a lush, rich sound."[5] The subjects knew that they needed to remember only the word in capitals. Later, experimenters administered a memory test for the words, with some hints. For *piano,* the hint was either "something heavy" or "something that makes music." The results showed that the subjects' memories were really good if the hint matched the way they had thought about *piano,* but poor if it didn't. That is, if the subjects read the moving men version of the sentence, hearing the cue "something that makes music" didn't help them remember *piano.* So it's not even enough to say, "You should think about meaning." You have to think about the right aspect of meaning.

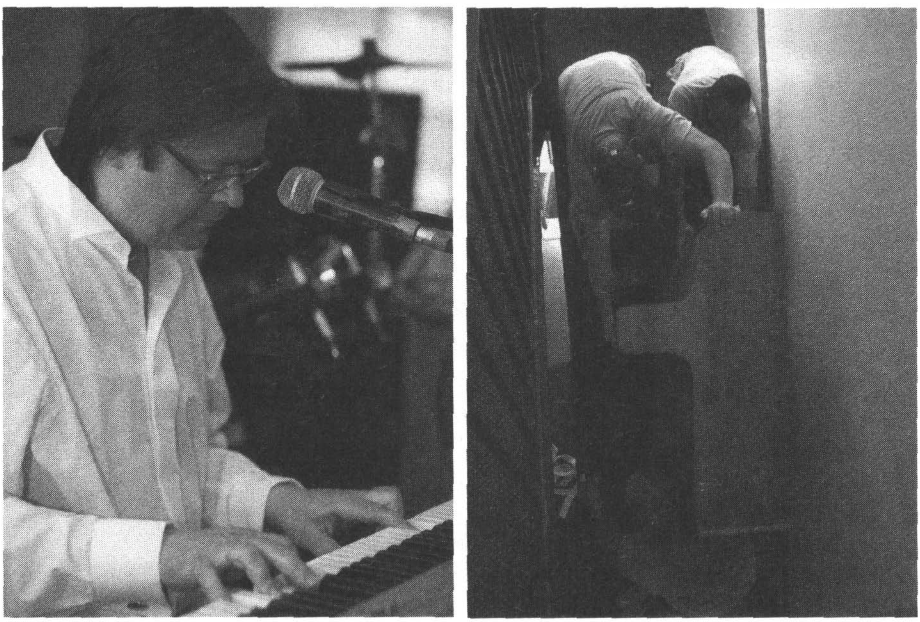

FIGURE 7: Two pictures of a piano, each emphasizing a different characteristic.

Let me summarize what I've said about learning so far. For material to be learned (that is, to end up in long-term memory), it must reside for some period in working memory—that is, a student must pay attention to it. Further, *how* the student thinks of the experience completely determines what will end up in long-term memory.

The obvious implication for teachers is that they must design lessons that will ensure that students are thinking about the meaning of the material. A striking example of an assignment that didn't work for this reason came from my nephew's sixth-grade teacher. He was to draw a plot diagram of a book he had recently finished. The point of the plot diagram was to get him to think about the story elements and how they related to one another. The teacher's goal, I believe, was to encourage her students to think of novels as having *structure,* but the teacher thought that it would be useful to integrate art into this project, so she asked her students to draw pictures to represent the plot elements. That meant that my nephew thought very little about the relation between different plot elements and a great deal about how to draw a good castle. My daughter had completed a similar assignment some years earlier, but her teacher had asked students to use words or phrases rather than pictures. I think that assignment more effectively fulfilled the intended goal because my daughter thought more about how ideas in the book were related.

Now you may be thinking, "OK, so cognitive psychologists can explain why students have to think about what material means—but I really already knew they should think about that. Can you tell me *how* to make sure that students think about meaning?" Glad you asked.

What Good Teachers Have in Common

If you read Chapter One, you can easily guess a common technique that I would *not* recommend for getting students to think about meaning: trying to make the subject matter relevant to the students' interests. I know that sounds odd, so let me elaborate.

Trying to make the material relevant to students' interests doesn't work. As I noted in Chapter One, content is seldom the decisive factor in whether or not our interest is maintained. For example, I love cognitive psychology, so you might think, "Well, to get Willingham to pay attention to this math problem, we'll wrap it up in a cognitive psychology example." But Willingham is quite capable of being bored by cognitive psychology, as has been proved repeatedly at professional conferences I've attended. Another problem with trying to use content to engage students is that it's sometimes very difficult to do and the whole enterprise comes off as artificial. How would a math instructor make algebra relevant to my sixteen-year-old daughter? With a "real-world" example using cell phone minutes? I just finished pointing out that any material has different aspects of meaning. If the instructor used a math problem with cell phone minutes, isn't there some chance that my daughter would think about cell phones rather than about the problem? And that thoughts about cell phones would lead to thoughts about the text message she received earlier, which would remind her to change her picture on her Facebook profile, which would make her think about the zit she has on her nose . . . ?

So if content won't do it, how about style? Students often refer to good teachers as those who "make the stuff interesting." It's not that the teacher relates the material to students' interests—rather, the teacher has a way of interacting with students that they find engaging. Let me give a few examples from my own experience with fellow college-level teachers who are consistently able to get students to think about meaning.

Teacher A is the comedian. She tells jokes frequently. She never misses an opportunity to use a silly example.

Teacher B is the den mother. She is very caring, very directive, and almost patronizing, but so warm that she gets away with it. Students call her "Mom" behind her back.

Teacher C is the storyteller. He illustrates almost everything with a story from his life. Class is slow paced and low key, and he is personally quiet and unassuming.

Teacher D is the showman. If he could set off fireworks inside, he would do it. The material he teaches does not lend itself easily to demonstrations, but he puts a good deal of time and energy into thinking up interesting applications, many of them involving devices he's made at home.

Each of these teachers is one to whom students refer as making boring material interesting, and each is able to get students to think about meaning. Each style works well for the person using it, although obviously not everyone would feel comfortable taking on some of these styles. It's a question of personality.

Style is what the students notice, but it is only a part of what makes these teachers so effective. College professors typically get written student evaluations of their teaching at the end of every course. Most schools have a form for students to fill out that includes such items as "The professor was respectful of student opinions," "The professor was an effective discussion leader," and so on, and students indicate whether or not they agree with each statement. Researchers have examined these sorts of surveys to figure out which professors get good ratings and why. One of the interesting findings is that most of the items are redundant. A two-item survey would be almost as useful as a thirty-item survey, because all of the questions really boil down to two: Does the professor seem like a nice person, and is the class well organized? (See Figure 8.) Although they don't realize they are doing so, students treat each of the thirty items as variants of one of these two questions.

Although K–12 students don't complete questionnaires about their teachers, we know that more or less the same thing is true for them. The emotional bond between

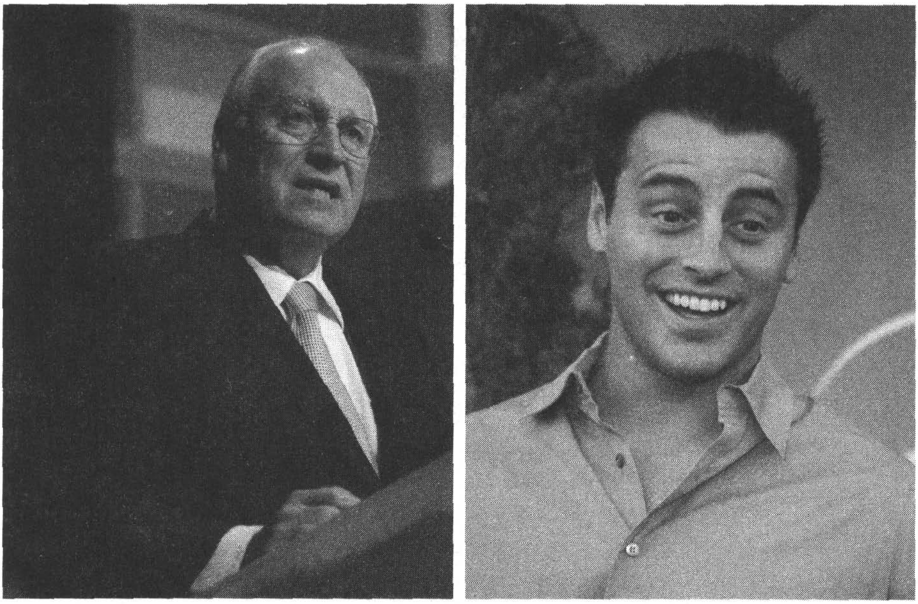

FIGURE 8: How would each of these men be as a teacher? Dick Cheney is smart but seems rather cold and forbidding. The character Joey Tribbiani from *Friends* (played by actor Matt LeBlanc) is warm and friendly but not terribly smart. Teachers need to be both well organized and approachable.

students and teacher—for better or worse—accounts for whether students learn. The brilliantly well-organized teacher whom fourth graders see as mean will not be very effective. But the funny teacher, or the gentle storytelling teacher, whose lessons are poorly organized won't be much good either. Effective teachers have both qualities. They are able to connect personally with students, and they organize the material in a way that makes it interesting and easy to understand.

That's my real point in presenting these different types of teachers. When we think of a good teacher, we tend to focus on personality and on the way the teacher presents himself or herself. But that's only half of good teaching. The jokes, the stories, and the warm manner all generate goodwill and get students to pay attention. But then how do we make sure they think about meaning? That is where the second property of being a good teacher comes in—organizing the ideas in a lesson plan in a coherent way so that students will understand and remember. Cognitive psychology cannot tell us how to be personable and likable to our students, but I can tell you about one set of principles that cognitive psychologists know about to help students think about the meaning of a lesson.

The Power of Stories

The human mind seems exquisitely tuned to understand and remember stories—so much so that psychologists sometimes refer to stories as "psychologically privileged," meaning that they are treated differently in memory than other types of material. I'm going to suggest that organizing a lesson plan like a story is an effective way to help students comprehend and remember. It also happens to be the organizing principle used by the four teachers I described. The way in which each of them related emotionally to their students was very different, but the way they got their students to think about the meaning of material was identical.

Before we can talk about how a story structure could apply to a classroom, we must go over what a story structure is. There is not universal agreement over what makes a story, but most sources point to the following four principles, often summarized as *the four Cs.* The first C is *causality,* which means that events are causally related to one another. For example, "I saw Jane; I left the house" is just a chronological telling of events. But if you read, "I saw Jane, my hopeless old love; I left the house," you would understand that the two events are linked causally. The second C is *conflict.* A story has a main character pursuing a goal, but he or she is unable to reach that goal. In *Star Wars* the main character is Luke Skywalker, and his goal is to deliver the stolen plans and help destroy the Death Star. Conflict occurs because there is an obstacle to the goal. If Luke didn't have a worthy adversary—Darth Vader—it would make for a rather short movie. In any story the protagonist must struggle to meet his goal. The third C is *complications.* If Luke simply hammered away for ninety minutes at his goal of delivering the plans, that would be rather dull. Complications are subproblems that arise from the main goal. Thus, if Luke wants to deliver the plans, he must first get off his home planet, Tatooine—but he has no transportation. That's a complication that leads to his meeting another major character, Han Solo, and leaving the planet amid a hail of gunfire—always a movie bonus. The final C is *character.* A good story is built around strong, interesting characters, and the key to those qualities is *action.* A skillful storyteller shows rather than tells the audience what a character is like. For example, the first time the *Star Wars* audience sees Princess Leia she is shooting at storm troopers. Hence we don't need to be told that she is brave and ready to take action.

If we're trying to communicate with others, using a story structure brings several important advantages. First, stories are easy to comprehend, because the audience knows the structure, which helps to interpret the action. For example, the audience knows that events don't happen randomly in stories. There must be a causal connection, so if the cause is not immediately apparent, the audience will think carefully about the previous action to try to connect it to present events. For example, at one point in *Star Wars* Luke, Chewbacca, and Han are hiding on an Empire ship. They need to get to another part of the ship, and Luke suggests putting handcuffs on Chewbacca. That suggestion is mildly puzzling because Luke and Chewbacca are allies. The audience must figure out that Luke intends to pretend that Chewbacca is a prisoner and that he and Han are guards. The audience will do that bit of mental work because they know there must be a reason for this puzzling action.

Second, stories are interesting. Reading researchers have conducted experiments in which people read lots of different types of material and rate each for how interesting it is. Stories are consistently rated as more interesting than other formats (for example, expository prose), even if the same information is presented. Stories may be interesting because they demand the kind of inferences I discussed in Chapter One. Recall that problems (such as crossword puzzles) are interesting if they are neither too difficult nor too easy. Stories demand these medium-difficulty inferences, as in the handcuff example just presented.

Formal work in laboratory settings has shown that people rate stories as less interesting if they include too much information, thus leaving no inferences for the listener to make. But formal research is hardly necessary to confirm this phenomenon. We all have one or two friends who kill every story they tell with too much information. (See Figure 9.) An acquaintance of mine recently spent ten minutes relating that the owner of her favorite Chinese restaurant, which she hadn't visited in a year because they no longer took checks, informed her that he would make an exception for her. Delivered in fifteen seconds with cheeky pride, this story would have been cute. But with the details packed in (and no inferences for me to make) over the course of ten full minutes, it was all I could do not to scream.

Third, stories are easy to remember. There are at least two contributing factors here. Because comprehending stories requires lots of medium-difficulty inferences, you must think about the story's meaning throughout. As described earlier in the chapter, thinking about meaning is excellent for memory because it is usually meaning that you want to remember. Your memory for stories is also aided by their causal structure. If you remember one part of the plot, it's a good guess that the next thing that happened was caused by what you remember. For example, if you're trying to remember what happened after Luke put handcuffs on Chewbacca, you'll be helped by remembering that they were on an Empire ship (hence the ruse), which might help you remember that they went to rescue Princess Leia from the detention area.

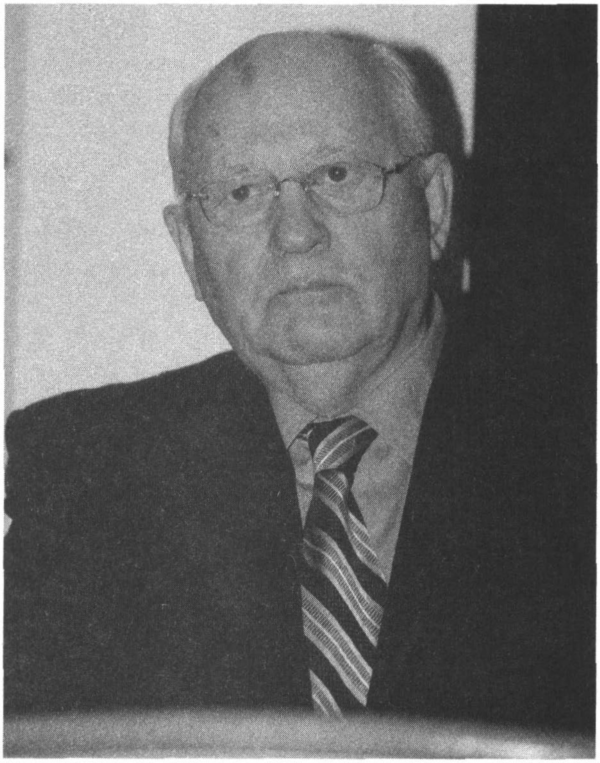

FIGURE 9: Former head of state of the Soviet Union Mikhail Gorbachev was well known to reporters for giving answers that were boring because they were exhaustive. At a 1990 question-and-answer session with a dozen members of the U.S. Congress, Gorbachev answered the first question (about the Soviet economy) with a twenty-eight-minute monologue that covered all aspects of property rights while senators looked "glazed" or "weary." Senator Robert Dole later remarked, "He does have long answers."[6]

Putting Story Structure to Work

Now, all this about movies has been a diverting interlude (at least I hope it has), but what does it have to do with the classroom? My intention here is not to suggest that you simply tell stories, although there's nothing wrong with doing so. Rather, I'm suggesting something one step removed from that. Structure your lessons the way stories are structured, using the four Cs: causality, conflict, complications, and character. This doesn't mean you must do most of the talking. Small group work or projects or any other method may be used. The story structure applies to the way you *organize* the material that you encourage your students to think about, not to the methods you use to teach the material.

In some cases, the way to structure a lesson plan as a story is rather obvious. For example, history can be viewed as a set of stories. Events are caused by other events; there is often conflict involved; and so on. Still, thinking carefully about the four Cs as you consider a lesson plan can be helpful. It might encourage you to think about a different perspective from which to tell the story. For example, suppose you are planning a lesson on Pearl Harbor. You might first think of the organization shown in Figure 10. It's chronological and it makes the United States the main character—that is, events are taken from the U.S. point of view. Your goal is to get students to think about three points: U.S. isolationism before Pearl Harbor, the attack, and the subsequent "Germany first" decision and the putting of the United States on a war footing.

Suppose, however, you thought of the four Cs when you were telling this story. From that perspective, the United States is not the strong character. Japan is, because she had the goal that propelled events forward—regional domination—and she had significant obstacles to this goal—she lacked natural resources and she was embroiled in a protracted war with China. This situation set up a subgoal: to sweep up the European colonies in the South Pacific. Meeting that goal would raise Japan's standing as a world power and

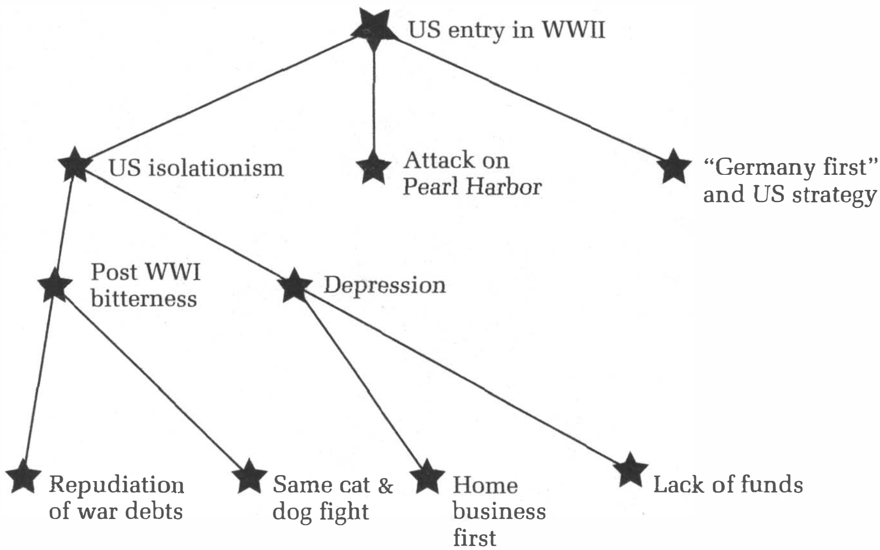

FIGURE 10: A tree diagram showing the typical structure of a lesson plan on Pearl Harbor. The organization is chronological.

FIGURE 11: Alternative organization for a lesson plan on Pearl Harbor. From a storytelling point of view, Japan is the strong character because she takes actions that move the story forward.

help her obtain crucial raw materials for finishing the war with China. But that subgoal brought with it another complication. The United States was the other major naval power in the Pacific. How was Japan to deal with that problem? Rather than plundering the European colonies and daring the United States to intervene across five thousand miles of ocean (which the United States probably would not have done), Japan chose to try to eliminate the threat in one surprise attack. If one seeks to organize a lesson plan as a story, the one in Figure 10 is less compelling than the one in Figure 11.

My suggestion to use the Japanese point of view of Pearl Harbor doesn't mean that the American point of view should be ignored or deemed less important. Indeed, I could imagine a teacher in the United States electing not to use this story structure precisely because it takes a Japanese point of view in a U.S. history class. My point here is that using a story structure may lead you to organize a lesson in ways that you hadn't considered before. And the story structure does bring cognitive advantages.

Using storytelling to teach history seems easy, but can you really use a story structure in a math class? Absolutely. Here's an example of how I introduced the concept of a Z-score—a common way to transform data—when I taught introductory statistics. Begin with the simplest and most familiar example of probability—the coin flip. Suppose I have a coin that I claim is loaded—it always comes up heads. To prove it to you, I flip the coin and it does indeed come up heads. Are you convinced? College students understand that the answer should be no because there is a fifty-fifty chance that a fair coin would have come up heads. How about one hundred heads in a row? Clearly the odds are really small that a fair coin will come up heads one hundred times in a row, so you'd conclude that the coin isn't fair.

That logic—how we decide whether a coin is fishy or fair—is used to evaluate the outcome of many, if not most, scientific experiments. When we see headlines in the newspaper saying "New drug for Alzheimer's found effective" or "Older drivers less safe than younger" or "Babies who watch videos have smaller vocabularies," these conclusions rest on the same logic as the coin flip. How?

Suppose we want to know whether an advertisement is effective. We ask two hundred people, "Does Pepsodent give you sex appeal?" One hundred of these people have seen an advertisement for Pepsodent and one hundred have not. We want to know if the percentage of people in the saw-the-ad group who say it gives you sex appeal is higher than the percentage in the didn't-see-the-ad group who say it gives you sex appeal. The problem here is just like the problem with the coin-flip example. The odds of the saw-the-ad group being higher are around 50 percent. One of the two groups *has* to be higher. (If they happened to tie, we'd assume that the ad didn't work.)

The logic for getting around this problem is the same as it was for the coin-flip example. For the coin flip, we judged one hundred heads in a row as a highly improbable event *assuming that the coin was fair.* The odds of a fair coin coming up heads one hundred times in a row are very small. So if we observe that event—one hundred heads in a row—we conclude that our assumption must have been wrong. It's *not* a fair coin. So the saw-the-ad group being higher than the other group may also not be improbable—but what if that group was *much* more likely to answer yes? Just as we judged that there was something funny about the coin, so too we should judge that there is something funny about people who have seen the ad—at least funny when it comes to answering our question.

Of course *funny* in this context means "improbable." In the case of the coin, we knew how to calculate the "funniness," or improbability, of events because we knew the number of possible outcomes (two) and the probability of each individual outcome (.5), so it was easy to calculate the odds of successive events, as shown in Table 1. But here's our next problem: How do we calculate the "funniness," or probability, of other types of events? How much worse does the vocabulary of kids who watched videos have to be compared to that of kids who didn't watch videos before we're prompted to say, "Hey, these two groups of kids are not equal. If they were equal, their vocabularies would be equal. But their vocabularies are *very* unequal."

All of this description of coins, advertisements, and experiments is really a prelude to the lesson. I'm trying to get students to understand and care about the goal of the lesson, which is to explain how we can determine the probability of an event occurring by

TABLE 1: The odds, out of ten tosses, of tossing a successively greater number of heads.

Number of Tosses	Approximate Probability of All Heads
1	.5
2	.25
3	.125
4	.063
5	.031
6	.016
7	.008
8	.004
9	.002
10	.001

chance. That is the conflict for this lesson. Our worthy adversary in pursuit of this goal is not Darth Vader but the fact that most events we care about are not like coin flips—they don't have a limited number of outcomes (heads or tails) for which we know the probabilities (50 percent). That's a complication, which we address with a particular type of graph called a histogram; but implementing this approach leads to a further complication: we need to calculate the area under the curve of the histogram, which is a complex computation. The problem is solved by the Z-score, which is the point of the lesson (Figure 12).

A couple of things are worth noticing. A good deal of time—often ten or fifteen minutes of a seventy-five-minute class—is spent setting up the goal, or to put it another way, persuading students that it's important to know how to determine the probability of a chance event. The material covered during this setup is only peripherally related to the lesson. Talking about coin flips and advertising campaigns doesn't have much to do with Z-scores. It's all about elucidating the central conflict of the story.

Spending a lot of time clarifying the conflict follows a formula for storytelling from, of all places, Hollywood. The central conflict in a Hollywood film starts about twenty minutes into the standard one-hundred-minute movie. The screenwriter uses that twenty minutes to acquaint you with the characters and their situation so that when the main conflict arises, you're already involved and you care what happens to the characters. A film may start with an action sequence, but that sequence is seldom related to what will be the main story line of the movie. James Bond movies often start with a chase scene, but it's always part of some other case, not the case that Bond will work on for the bulk of the movie. The conflict for that case is introduced about twenty minutes into the film.

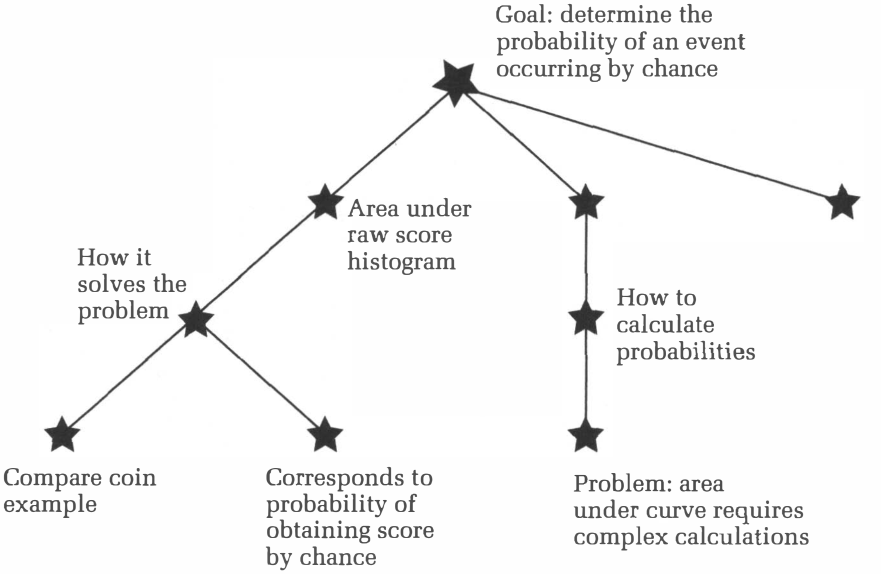

FIGURE 12: Part of the organizational scheme for a lesson plan on the Z score transformation for a statistics class.

When it comes to teaching, I think of it this way: The material I want students to learn is actually the answer to a question. *On its own, the answer is almost never interesting.* But if you know the question, the answer may be quite interesting. That's why making the question clear is so important. But I sometimes feel that we, as teachers, are so focused on getting to the answer, we spend insufficient time making sure that students understand the question and appreciate its significance.

Let me close this section by emphasizing again that there are many ways in which one can be a good teacher. I don't mean to imply that, according to cognitive science, every teacher should be using a story structure to shape his or her lesson plans. It's just one way that we can help ensure that students think about meaning. I am implying—well, no, I'm stating—that every teacher should get his or her students to think about the meaning of material—except sometimes, which is the subject of the next section.

But What If There Is No Meaning?

This chapter began by posing the question, *How can we get students to remember something?* The answer from cognitive science is straightforward: get them to think about what it means. In the previous section I suggested one method—story structure—for getting students to think about meaning.

It's fair to ask, however, whether there is material that students must learn that is pretty darn close to meaningless. For example, how can you emphasize meaning when students are learning the odd spelling of *Wednesday,* or that *enfranchise* means to give voting rights, or that *travailler* is the French verb for *work?* Some material just doesn't seem to have much meaning. Such material seems especially prevalent when one is entering a new field or domain of knowledge. A chemistry teacher might want students to learn in order the symbols for a few elements of the periodic table—but how can students think of the symbols H, He, Li, Be, B, C, N, O, and F in a deep, meaningful way when they don't know any chemistry?

Memorizing meaningless material is commonly called *rote memorization.* I will say more about what rote memory really is in Chapter Four, but for the moment let's just acknowledge that a student who has memorized the first nine elements of the periodic table has little or no idea why she has done so or what the ordering might mean. There are times when a teacher may deem it important for a student to have such knowledge ready in long-term memory as a stepping-stone to understanding something deeper. How can a teacher help the student get that material into long-term memory?

There is a group of memory tricks, commonly called *mnemonics,* that help people memorize material when it is not meaningful. Some examples are listed in Table 2.

I'm not a big fan of the peg-word and method-of-loci methods because they are hard to use for different sets of material. If I use my mental walk (back porch, dying pear tree, gravel driveway, and so on) to learn some elements of the periodic table, can I use the same walk to learn the conjugations for some French verbs? The problem is

TABLE 2: Common mnemonic methods. Mnemonics help you to memorize meaningless material.

Mnemonic	How It Works	Example
Peg word	Memorize a series of *peg words* by using a rhyme—for example, one is a bun, two is a shoe, three is a tree, and so on. Then memorize new material by associating it via visual imagery with the pegs.	To learn the list *radio, shell, nurse* you might imagine a radio sandwiched in a bun, a shoe on a beach with a conch in it, and a tree growing nurses' hats like fruit.
Method of loci	Memorize a series of locations on a familiar walk— for example, the back porch of your house, a dying pear tree, your gravel driveway, and so on. Then visualize new material at each "station" of the walk.	To learn the list *radio, shell, nurse* you might visualize a radio hanging by its cord on the banister of your back porch, someone grinding shells to use as fertilizer to revitalize the dying tree, and a nurse shoveling fresh gravel onto your driveway.
Link method	Visualize each of the items connected to one another in some way.	To learn the list *radio, shell, nurse* you might imagine a nurse listening intently to a radio while wearing large conch shells on her feet instead of shoes.
Acronym method	Create an acronym for the to-be-remembered words, then remember the acronym.	To learn the list *radio, shell, nurse* you might memorize the word RAiSiN using the capitalized letters as cues for the first letter of each word you are to remember.
First letter method	Similar to the acronym method, this method has you think of a phrase, the first letter of which corresponds to the first letter of the to-be-remembered material.	To learn the list *radio, shell,* nurse you could memorize the phrase "Roses smell nasty," then use the first letter of each word as a cue for the words on the list.
Songs	Think of a familiar tune to which you can sing the words.	To learn the list *radio, shell, nurse* you could sing the words to the tune of "Happy Birthday to You."

that there might be interference between the two lists; when I get to the gravel driveway, I get confused about what's there because I've associated two things with it.

The other methods are more flexible because students can create a unique mnemonic for each thing they learn. The acronym method and the first-letter method are effective, but students do need to have some familiarity with the material to be learned. I always think of the acronym HOMES when I'm trying to remember the names of the great lakes. If I didn't already know the names, these first-letter cues wouldn't do me much good, but the first letter of each lake's name pushes me over the edge from tip-of-the-tongue to ready recall. The first-letter method works in much the same way, and has the same limitation.

Setting to-be-learned information to music or chanting it to a rhythm also works quite well. Most of us learned the letters of the alphabet by singing the ABC song, and I've seen the state capitals set to the music of the "Battle Hymn of the Republic." Music and rhythm do make words remarkably memorable, and the song doesn't have to be particularly melodic. I can still remember the character Coach from the television show *Cheers* studying for a geography exam by singing (to the tune of "When the Saints Go Marching In"),

> Albania! Albania! You border on the Adriatic.
> Your land is mostly mountainous, and your main export is chrome.

The difficulty with songs is that they are more difficult to generate than the other mnemonic devices.

Why do mnemonics work? Primarily by giving you cues. The acronym ROY G. BIV gives you the first letter of each color in the spectrum of visible light. The first letter is quite a good cue to memory. As I discuss in the next chapter, memory works on the basis of cues. If you don't know anything about a topic, or if the things you're trying to remember are confusing because they are arbitrary (there's nothing about red that makes it obvious that its wavelength is longer than green), mnemonics help because they impose some order on the material.

Let me summarize what I've said in this chapter. If we agree that background knowledge is important, then we must think carefully about how students can acquire that background knowledge—that is, how learning works. Learning is influenced by many factors, but one factor trumps the others: students remember what they think about. That principle highlights the importance of getting students to think about the right thing at the right time. We usually want students to understand what things ***mean***, which sets the agenda for a lesson plan. How can we ensure that students think about meaning? I offered one suggestion, which is to use the structure of a story. Stories are easily comprehended and remembered, and they are interesting; but one can't get students to *think* about meaning if the material *has* no meaning. In that case, it may be appropriate to use a mnemonic device.

Implications for the Classroom

Thinking about meaning helps memory. How can teachers ensure that students think about meaning in the classroom? Here are some practical suggestions.

Review Each Lesson Plan in Terms of What the Student Is Likely to Think About

This sentence may represent the most general and useful idea that cognitive psychology can offer teachers. The most important thing about schooling is what students will remember after the school day is over, and there is a direct relationship between what they think during the day and their later memory. So it's a useful double-check for every lesson plan to try to anticipate what the lesson will *actually* make students think about (rather than what you hope it will make them think about). Doing so may make it clear that students are unlikely to get what the teacher intended out of the lesson.

For example, I once observed a high school social studies class work in groups of three on projects about the Spanish Civil War. Each group was to examine a different aspect of the conflict (for example, compare it to the U.S. Civil War, or consider its impact on today's Spain) and then teach the remainder of the class what they had learned, using the method of their choice. The teacher took students to a computer laboratory to do research on the Internet. (They also used the library.) The students in one group noticed that PowerPoint was loaded on the computers, and they were very enthusiastic about using it to teach their bit to the other groups. The teacher was impressed by their initiative and gave his permission. Soon all of the groups were using PowerPoint. Many students had some familiarity with the basics of the program, so it could have been used effectively. The problem was that the students changed the assignment from "learn about the Spanish Civil War" to "learn esoteric features of PowerPoint." There was still a lot of enthusiasm in the room, but it was directed toward using animations, integrating videos, finding unusual fonts, and so on. At that point the teacher felt it was far too late to ask all of the groups to switch, so he spent much of the rest of the week badgering students to be sure their presentation had content, not just flash.

This story illustrates one of the reasons that experienced teachers are so good. This teacher clearly didn't let students use PowerPoint the next year, or possibly he thought of a way to keep them on task. Before you have accumulated these experiences, the next best thing is to think carefully about how your students will react to an assignment, and what it will make them think about.

Think Carefully About Attention Grabbers

Almost every teacher I have met likes, at least on occasion, to start class with an attention grabber. If you hook students early in the lesson, they should be curious to know what is behind whatever surprised or awed them. But attention grabbers may not always work. Here's a conversation I had with my oldest daughter when she was in sixth grade.

Dad: What did you do in school today?

Rebecca: We had a guest in science. He taught us about chemicals.

Dad: Oh yeah? What did you learn about chemicals?

Rebecca: He had this glass? That looked like water? But when he put this little metal thingy in it, it boiled. It was so cool. We all screamed.

Dad: Uh-huh. Why did he show you that?

Rebecca: I don't know.

The guest surely planned this demonstration to pique the class's interest, and that goal was met. I'm willing to bet that the guest followed the demonstration with an age-appropriate explanation of the phenomenon but that information was not retained. Rebecca didn't remember it because she was still thinking about how cool the demonstration was. You remember what you think about.

Another teacher once told me she wore a toga to class on the first day she began a unit on ancient Rome. I am sure that got her students' attention. I am also sure it continued to get their attention—that is, to distract them—once the teacher was ready for them to think about something else.

Here's one more example. A guest in a biology class asked the students to think of the very first thing they had ever seen. The students mulled that question over and generated such guesses as "the doctor who pulled me out," "Mom," and so forth. The guest then said, "Actually, the first thing each of you saw was the same. It was pinkish, diffuse light coming through your mother's belly. Today we're going to talk about how that first experience affected how your visual system developed, and how it continues to influence the way you see today." I love that example because it grabbed the students' attention and left them eager to hear more about the subject of the lesson.

As I alluded to earlier in the chapter, I think it is very useful to use the beginning of class to build student interest in the material, or as I put it, to develop the conflict. You might consider, however, whether the beginning of the class is really when they need an attention grabber. In my experience, the transition from one subject to another (or for older students, from one classroom and teacher to another) is enough to buy at least a few minutes of attention from students. It's usually the middle of the lesson that needs a little drama to draw students back from whatever reverie they might be in. But regardless of when it's used, think hard about how you will draw a connection between the attention grabber and the point it's designed to make. Will students understand the connection, and will they be able to set aside the excitement of the attention grabber and move on? If not, is there a way to change the attention grabber to help students make that transition? Perhaps the toga could be worn over street clothes and removed after the first few minutes of class. Perhaps the "metal thingy" demonstration would have been better *after* the basic principle was explained and students were prompted to predict what might happen.

Use Discovery Learning with Care

In discovery learning students learn by exploring objects, discussing problems with classmates, designing experiments, or any of a number of other techniques that use student inquiry rather than have the teacher tell students things. Indeed, the teacher ideally serves more as a resource than as the director of the class. Discovery learning has much to recommend it, especially when it comes to the level of student engagement. If students have a strong voice in deciding which problems they want to work on, they will likely be engaged in the problems they select, and will likely think deeply about the material, with attendant benefits. An important downside, however, is that what students will think about is less predictable. If students are left to explore ideas on their own, they may well explore mental paths that are not profitable. If memory is the residue of thought, then students will remember incorrect "discoveries" as much as they will remember the correct ones.

Now this doesn't mean that discovery learning should never be used, but it does suggest a principle for when to use it. Discovery learning is probably most useful when the environment gives prompt feedback about whether the student is thinking about a problem in the right way. One of the best examples of discovery learning is when kids learn to use a computer, whether they are learning an operating system, a complex game, or a Web application. Kids show wonderful ingenuity and daring under these circumstances. They are not afraid to try new things, and they shrug off failure. They learn by discovery! Note, however, that computer applications have an important property: when you make a mistake, it is immediately obvious. The computer does something other than what you intended. This immediate feedback makes for a wonderful environment in which "messing around" can pay off. (Other environments aren't like that. Imagine a student left to "mess around" with frog dissection in a biology class.) If the teacher does not direct a lesson to provide constraints on the mental paths that students will explore, the environment itself can do so effectively in a discovery learning context, and that will help memory.

Design Assignments So That Students Will Unavoidably Think About Meaning

If the goal of a lesson plan is to get students to think about the meaning of some material, then it's pretty clear that the best approach is one in which thinking about meaning is unavoidable. One of the things that has always amazed me as a memory researcher is the degree to which people do not know how their own memory system works. It doesn't do any good to tell people, "Hey, I'm going to test your memory for this list of words later," because people don't know what to do to make the words memorable. But if you give people a simple task in which they *must* think of the meaning—for example, rating how much they like each word—they will remember the words quite well.

This idea can be used in the classroom as well as in the laboratory. At the start of this chapter I said that asking fourth graders to bake biscuits was not a good way to get them to appreciate what life on the Underground Railroad was like because they spend too much time thinking about measuring flour and milk. The goal was to get students thinking about the experience of runaway slaves. So a more effective lesson would be to lead students to consider that experience by, for example, asking them where they supposed runaway slaves obtained food, how they were able to prepare it, how they were able to pay for it, and so forth.

Don't Be Afraid to Use Mnemonics

Many teachers I have met shudder at the use of mnemonics. They conjure up images of nineteenth-century schoolrooms with children chanting rhymes of the state capitals. But as bad as a classroom would be if a teacher used *only* mnemonics, they do have their time and place, and I don't think teachers should have this instructional technique taken away from them.

When is it appropriate to ask students to memorize something before it has much meaning? Probably not often, but there will be times when a teacher feels that some material—meaningless though it may be now—must be learned for the student to move forward. Typical examples would be learning letter-sound associations prior to reading, and learning vocabulary in both their native language and foreign languages.

It might also be appropriate to memorize some material using mnemonics in parallel with other work that emphasizes meaning. When I was in elementary school, I was not required to memorize the multiplication table. Instead I practiced using different materials and techniques that emphasized what multiplication actually means. These techniques were effective, and I readily grasped the concept. But by about fifth grade, not knowing the multiplication table by heart really slowed me down because the new things I was trying to learn had multiplication embedded in them. So every time I saw 8×7 within a problem I had to stop and figure out the product. In the sixth grade I moved to a new school, where my teacher quickly figured out what was going on and made me memorize the multiplication table. It made math a lot easier for me, although it took a few weeks before I would admit it.

Try Organizing a Lesson Plan Around the Conflict

There is a conflict in almost any lesson plan, if you look for it. This is another way of saying that the material we want students to know is the answer to a question—and the question is the conflict. The advantage of being very clear about the conflict is that it yields a natural progression for topics. In a movie, trying to resolve a conflict leads to new complications. That's often true of school material too.

Start with the material you want your students to learn, and think backward to the intellectual question it poses. For example, in a science class you might want sixth graders to know the models of the atom that were competing at the turn of the twentieth century. These are the answers. What is the question? In this story, the goal is to understand the nature of matter. The obstacle is that the results of different experiments appear to conflict with one another. Each new model that is proposed (Rutherford, cloud, Bohr) seems to resolve the conflict but then generates a new complication—that is, experiments to test the model seem to conflict with other experiments. If this organization seems useful to you, you might spend a good bit of time thinking about how to illustrate and explain to students the question, "What is the nature of matter?" Why should that question interest sixth graders?

As I've emphasized, structuring a lesson plan around conflict can be a real aid to student learning. Another feature I like is that, if you succeed, you are engaging students with the actual substance of the discipline. I've always been bothered by the advice "make it relevant to the students," for two reasons. First, it often feels to me that it doesn't apply. Is the Epic of Gilgamesh relevant to students in a way they can understand right now? Is trigonometry? Making these topics relevant to students' lives will be a strain, and students will probably think it's phony. Second, if I can't convince students that some material is relevant, does that mean I shouldn't teach it? If I'm continually trying to build bridges between students' daily lives and their school subjects, the students may get the message that school is always about them, whereas I think there is value, interest, and beauty in learning about things that don't have much to do with me. I'm not saying it never makes sense to talk about things students are interested in. What I'm suggesting is that student interests should not be the main driving force of lesson planning. Rather, they might be used as initial points of contact that help students understand the main ideas you want them to consider, rather than as the reason or motivation for them to consider these ideas.

In the previous chapter I argued that students must have background knowledge in order to think critically. In this chapter I discussed how memory works, in the hope that by understanding this we can maximize the likelihood that students will learn this background knowledge; much of the answer to how we can do this was concerned with thinking about meaning. But what if students don't understand the meaning? In the next chapter I discuss why it is hard for students to comprehend the meaning of complex material, and what you can do to help.

Note

* I made up this statistic.

Bibliography

Less Technical

Druxman, M. B. (1997). The art of storytelling: How to write a story . . . any story. Westlake Village, CA: Center Press. If you are interested in learning more about how stories are structured, this is a readable instruction manual.

Schacter, D. L. (2002). The seven sins of memory: How the mind forgets and remembers. Boston: Houghton Mifflin. A very readable account of why we remember and forget, with lots of examples that the reader can relate to, as well as descriptive studies of people with brain damage.

More Technical

Britton, B. K., Graesser, A. C., Glynn, S. M., Hamilton, T., & Penland, M. (1983). Use of cognitive capacity in reading: Effects of some content features of text. *Discourse Processes, 6,* 39–57. A study showing that people find stories more interesting than other types of text, even when they contain similar information.

Kim, S-i. (1999). Causal bridging inference: A cause of story interestingness. *British Journal of Psychology, 90*, 57–71. In this study the experimenter varied the difficulty of the inference that readers had to make to understand the text, and found that texts were rated as most interesting when the inferences were of medium-level difficulty.

Markman, A. B. (2002). Knowledge representation. In H. D. Pashler & D. L. Medin (Eds.), *Steven's handbook of experimental psychology*, Vol. 2: *Memory and cognitive processes.* (3rd ed., pp. 165–208). Hoboken, NJ: Wiley. A thorough treatment of how memories are represented in the mind, and of what representation actually means.

Meredith, G. M. (1969). Dimensions of faculty-course evaluation. *Journal of Psychology: Interdisciplinary and Applied, 73,* 27–32. An article showing that college students' attitudes toward professors are determined mostly by whether the professor is organized and seems nice. Not every study on this topic breaks it down in exactly this way, but this is the typical result.