ISSUE BRIEF #1, March 2016

www.KnowledgeMattersCampaign.org

Knowledge Matters

Restoring Wonder and Excitement to the Classroom

Top Three Takeaways

1. Critical thinking is not a set of skills and strategies that can be directly taught, practiced, and applied to any topic.

2. Students need deep knowledge of a subject in order to think creatively or critically about it.

3. There are no shortcuts to expert thinking. To "think like a scientist," a student must know the facts, concepts, and procedures that a scientist knows.

Daniel T. Willingham is a professor of cognitive psychology at the University of Virginia and author of several books, including the bestseller Why Don't Students Like School?

Knowledge and Practice: The Real Keys to Critical Thinking

By Daniel T. Willingham

Virtually everyone would agree that a primary, yet insufficiently met, goal of schooling is to enable students to think critically. In layperson's terms, critical thinking consists of seeing both sides of an issue, being open to new evidence that disconfirms your ideas, reasoning dispassionately, demanding that claims be backed by evidence, deducing and inferring conclusions from available facts, solving problems, and so forth. In addition, there are specific types of critical thinking that are characteristic of different subject matter: That's what we mean when we refer to "thinking like a scientist" or "thinking like a historian." This proper and commonsensical goal has very often been translated into calls to teach "critical thinking skills" and "higher-order thinking skills"—and into generic calls for teaching students to make better judgments, reason more logically, and so forth.

These calls are not new. In 1983, *A Nation At Risk*, a report by the National Commission on Excellence in Education, found that many 17-year-olds did not possess the "'higher-order' intellectual skills" this country needed. It claimed that nearly 40 percent could not draw inferences from written material and only one-fifth could write a persuasive essay. Following the release of *A Nation At Risk*, programs designed to teach students to think critically across the curriculum became extremely popular. By 1990, most states had initiatives designed to encourage educators to teach critical thinking, and one of the most widely used programs, Tactics for Thinking, sold 70,000 teacher guides.[1] But, for reasons I'll explain, the programs were not very effective; we still lament students' lack of critical thinking.

After more than 20 years of lamentation, exhortation, and little improvement, maybe it's time to ask a fundamental question: Can critical thinking actually be taught? Decades of cognitive research point to a disappointing answer: *not really*. People who have sought to teach critical thinking have assumed that it is a skill, like riding a bicycle, and that, like other skills, once you learn it, you can apply it in any situation. Research from cognitive science shows that thinking is not that sort of skill. The processes of thinking are intertwined with the content of thought—that is, domain knowledge. Thus, if you remind a student to "look at an issue from multiple perspectives" often enough, he will learn that he ought to do so, but if he doesn't know much about an issue, he can't think about it from multiple perspectives.

Critical thinking is not a set of skills that can be deployed at any time, in any context. It is a type of thought that even 3-year-olds can engage in—and even trained scientists can fail in. And it very much depends on domain knowledge and practice.

The processes of thinking are intertwined with the content of thought—that is, domain knowledge.

When one is very familiar with a problem's deep structure, knowledge about how to solve it transfers well.

Broad Knowledge for Comprehension, Deep Knowledge for Analysis

Anything you hear or read is automatically interpreted in light of what you already know about similar subjects. For example, suppose you read these two sentences: "After years of pressure from the film and television industry, the President has filed a formal complaint with China over what U.S. firms say is copyright infringement. These firms assert that the Chinese government sets stringent trade restrictions for U.S. entertainment products, even as it turns a blind eye to Chinese companies that copy American movies and television shows and sell them on the black market." Your broad background knowledge not only allows you to comprehend the sentences, it also has a powerful effect as you continue to read because it narrows the interpretations of new text that you will entertain. For example, if you later read the word "piracy," you would not think of eye-patched swabbies shouting "shiver me timbers!" The cognitive system gambles that incoming information will be related to what you've just been thinking about. Thus, it significantly narrows the scope of possible interpretations of words, sentences, and ideas. The benefit is that comprehension proceeds faster and more smoothly; the cost—as I explain below—is that the deep structure of a problem is harder to recognize.

Imagine a seventh-grade math class immersed in word problems. How is it that students will be able to answer one problem, but not the next, even though mathematically both word problems are the same, that is, they rely on the same mathematical knowledge? Typically, the students are focusing on the scenario that the word problem describes—its surface structure—instead of on the mathematics required to solve it—its deep structure. So even though students have been taught how to solve a particular type of word problem, when the teacher or textbook changes the scenario, students still struggle to apply the solution because they don't recognize that the problems are mathematically the same.

If knowledge of how to solve a problem never transferred to problems with new surface structures, schooling would be inefficient or even futile—but of course, such transfer does occur. When and why is complex,[2] but two factors are especially relevant for educators: familiarity with a problem's deep structure and the knowledge that one should look for a deep structure. I'll address each in turn.

When one is very familiar with a problem's deep structure, knowledge about how to solve it transfers well. That familiarity can come from long-term, repeated experience with one problem, or with various manifestations of one type of problem (i.e., many problems that have different surface structures, but the same deep structure). After repeated exposure to either or both, the subject simply perceives the deep structure as part of the problem description. Here's an example:

> A treasure hunter is going to explore a cave up on a hill near a beach. He suspected there might be many paths inside the cave so he was afraid he might get lost. Obviously, he did not have a map of the cave; all he had with him were some common items such as a flashlight and a bag. What could he do to make sure he did not get lost trying to get back out of the cave later?

The solution is to carry some sand with you in the bag, and leave a trail as you go, so you can trace your path back when you're ready to leave the cave. About 75 percent of American college students thought of this solution—but only 25 percent of Chinese students solved it.[3] The experimenters suggested that Americans solved it because most grew up hearing the story of Hansel and Gretel, which includes the idea of leaving a trail as you travel to an unknown place in order to find your way back. When the experimenters gave subjects another puzzle that shared a deep structure with a common Chinese folk tale, the percentage of solvers from each culture reversed.

It takes a good deal of practice with a problem type before students know it well enough to immediately recognize its deep structure, irrespective of the surface structure, as Americans did for the Hansel and Gretel problem. American subjects didn't think of the problem in terms of sand, caves, and treasure; they thought of it in terms of finding something with which to leave a trail. The deep structure of the problem is so well represented in their memory that they immediately saw that structure when they read the problem.

Now let's turn to the second factor that aids in transfer despite distracting differences in surface structure—knowing to look for a deep structure. Consider what would happen if I said to an American student who was struggling with the cave problem, "this is similar to Hansel and Gretel." The student would understand that the problems must share a deep structure and would try to figure out what it is. Students can do something similar without the hint. A student might think "I'm seeing this problem in a math class, so there must be a math formula that will solve this problem." Then he could scan his memory (or textbook) for candidates, and see if one of them helps. This is an example of what psychologists call metacognition, or regulating one's thoughts. In the introduction, I mentioned that you can teach students maxims about how they ought to think. Cognitive scientists refer to these maxims as metacognitive strategies. They are little chunks of knowledge—like "look for a problem's deep structure" or "consider both sides of an issue"—that students can learn and then use to steer their thoughts in more productive directions.

Helping students become better at regulating their thoughts was one of the goals of the critical thinking programs that were popular 20 years ago, but not very effective. Their modest benefit is likely due to teaching students to effectively use metacognitive strategies. Students learn to avoid biases that most of us are prey to when we think, such as settling on the first conclusion that seems reasonable, only seeking evidence that confirms one's beliefs, ignoring countervailing evidence, overconfidence, and others.[4] Thus, a student who has been encouraged many times to see both sides of an issue, for example, is probably more likely to spontaneously think "I should look at both sides of this issue" when working on a problem.

Unfortunately, metacognitive strategies can only take you so far. Although they suggest what you ought to do, they don't provide the knowledge necessary to implement the strategy. For example, you may know that you ought not accept the first reasonable-sounding solution to a problem, but that doesn't mean you know how to come up with alternative solutions or weigh how reasonable each one is. That requires domain knowledge and practice in putting that knowledge to work.

You may know that you ought not accept the first reasonable-sounding solution to a problem, but that doesn't mean you know how to come up with alternative solutions.

Knowledge and skills are intertwined, so it is unfortunate that in education we tend to talk about them as separate. As Andrew Rotherham and I wrote:[5]

> If you believe that skills and knowledge are separate, you are likely to draw two incorrect conclusions. First, because content is readily available in many locations but thinking skills reside in the learner's brain, it would seem clear that—if we must choose between them—skills are essential, whereas content is merely desirable. Second, if skills are independent of content, we could reasonably conclude that we can develop these skills through the use of any content. For example, if students can learn how to think critically about science in the context of any scientific material, a teacher should select content that will engage students (for instance, the chemistry of candy), even if that content is not central to the field. But all content is not equally important to mathematics, or to science, or to literature. To think critically, students need the knowledge that is central to the domain.

Since critical thinking relies so heavily on domain knowledge, educators may wonder if thinking critically in a particular domain is easier to learn. The quick answer is yes, it's a little easier. To understand why, let's focus on one domain, science, and examine the development of scientific thinking.

Thinking Like a Scientist

Experts in teaching science recommend that scientific reasoning be taught in the context of rich subject-matter knowledge. A committee of prominent science educators brought together by the National Research Council[6] put it plainly: "Teaching content alone is not likely to lead to proficiency in science, nor is engaging in inquiry experiences devoid of meaningful science content."

The committee drew this conclusion based on evidence that background knowledge is necessary to engage in scientific thinking. For example, consider devising a research hypothesis. One could generate multiple hypotheses for any given situation. Suppose you know that car A gets better gas mileage than car B and you'd like to know why. There are many differences between the cars, so which will you investigate first? Engine size? Tire pressure? Paint color? A key determinant of the hypothesis you select is plausibility. One's judgment about the plausibility of a factor is based on one's knowledge of the domain.

Other data indicate that familiarity with the domain makes it easier to juggle different factors simultaneously, which in turn allows you to construct experiments that simultaneously control for more factors. For example, in one experiment,[7] eighth-graders completed two tasks. In one, they were to manipulate conditions in a computer simulation to keep imaginary creatures alive. In the other, they were told that they had been hired by a swimming pool company to evaluate how the surface area of swimming pools was related to the cooling rate of its water. Students were more adept at designing experiments for the first task than the second, which the researchers interpreted as being due to students' familiarity with the relevant variables. Students are used to thinking about factors that might influence creatures' health (e.g., food, predators), but have less experience working with factors that might influence water temperature (e.g., volume, surface area). Hence, it is not the case that "controlling variables in an experiment" is a pure process that is not affected by subjects' knowledge of those variables.

Subjects who started with more and better integrated knowledge planned more informative experiments and made better use of experimental outcomes.

Prior knowledge and beliefs not only influence which hypotheses one chooses to test, they influence how one interprets data from an experiment. In one experiment,[8] undergraduates were evaluated for their knowledge of electrical circuits. Then they participated in three weekly, 1.5-hour sessions during which they designed and conducted experiments using a computer simulation of circuitry, with the goal of learning how circuitry works. The results showed a strong relationship between subjects' initial knowledge and how much subjects learned in future sessions, in part due to how the subjects interpreted the data from the experiments they had conducted. Subjects who started with more and better integrated knowledge planned more informative experiments and made better use of experimental outcomes.

Other studies have found similar results and have found that anomalous, or unexpected, outcomes may be particularly important in creating new knowledge—and particularly dependent upon prior knowledge.[9] Data that seem odd because they don't fit one's mental model of the phenomenon under investigation are highly informative. They tell you that your understanding is incomplete and they guide the development of new hypotheses. But you could only recognize the outcome of an experiment as anomalous if you had some expectation of how it would turn out. And that expectation would be based on domain knowledge, as would your ability to create a new hypothesis that takes the anomalous outcome into account.

The idea that scientific thinking must be taught hand in hand with scientific content is further supported by research on scientific problem solving; that is, when students calculate an answer to a textbook-like problem, rather than design their own experiment. A meta-analysis[10] of 40 experiments investigating methods for teaching scientific problem solving showed that effective approaches were those that focused on building complex, integrated knowledge bases as part of problem solving, for example by including exercises like concept mapping. Ineffective approaches focused exclusively on the strategies to be used in problem solving while ignoring the knowledge necessary for the solution.

What do all these studies boil down to? First, critical thinking (as well as scientific thinking and other domain-based thinking) is not a skill. There is not a set of critical thinking skills that can be acquired and deployed regardless of context. Second, there are metacognitive strategies that, once learned, make critical thinking more likely. Third, the ability to think critically—to actually do what the metacognitive strategies call for—depends on domain knowledge and practice. For teachers, the situation is not hopeless, but no one should underestimate the difficulty of teaching students to think critically.

The ability to think critically—to actually do what the metacognitive strategies call for—depends on domain knowledge and practice.

When schools use a knowledge-rich curriculum, students have many incidental learning opportunities.

The job of teaching content is doubly serious for teachers in preschool and early elementary classrooms.

Classroom Application: Building Students' Knowledge

One sometimes hears that the real goal of education is "learning to learn." As the proverb says, "Give a man a fish, and he will eat for a day; teach a man to fish, and he will eat for a lifetime." Better to teach students how to learn facts on their own, rather than teach them facts. The idea sounds appealing, but if it's coupled with the idea that teachers should emphasize cognitive processes (like comprehension and reasoning strategies), and place less emphasis on content, then it's counterproductive.

Many of the cognitive skills we want our students to develop—especially reading with understanding and successfully analyzing problems—are intimately intertwined with knowledge of content. When students learn facts they are not just acquiring grist for the mill—they are enabling the mill to operate more effectively. Background knowledge is absolutely integral to effectively deploying important cognitive processes. What does this mean for teachers?

1. Facts should be meaningful. "Fact learning" should not be understood as "rote memorization." The importance of knowledge to cognition does not mean that teachers should assign lists of facts for their students to memorize. Facts are useful only if they are meaningfully connected to other bits of knowledge. So, fact learning should be thought of as the kind of learning that results from, for example, reading a richly detailed biography—not a barren timeline of a person's life. Teachers should include opportunities for students to learn new material about the world and connect it to prior knowledge wherever possible. Mindless drilling is not an effective vehicle for building students' store of knowledge.

2. Knowledge acquisition can be incidental. Every fact that students learn need not be explicitly taught—students can learn facts incidentally. Incidental learning refers to learning that occurs when you are not specifically trying to learn. Much of what you know stuck in your memory not as a result of your consciously trying to remember it, but as a byproduct of thinking about it, such as when you reflect on a novel word that someone used in conversation or are fascinated by a new fact. When schools use a knowledge-rich curriculum, students have many incidental learning opportunities as they are immersed in meaningful, connected facts throughout the day. Teachers can also look for extra opportunities to provide incidental learning opportunities for their students, for example, by using a vocabulary word that the students likely do not know, but the meaning of which is deducible from the context of the sentence.

3. Knowledge learning should start early. Building a store of knowledge works like compound interest—it grows exponentially. For that reason, the earlier students add to their database of knowledge, the better. This process begins at home, long before children attend school. (Note that virtually all learning before children start school is incidental.) All teachers should take the job of teaching content to students seriously, but this job is doubly serious for teachers in preschool and early elementary classrooms. Because of the exponential learning rate, once children fall behind their peers, it becomes increasingly difficult to catch up. These young children can learn little, if any, material via reading, so they must learn by listening to fiction and nonfiction books read aloud, by watching demonstrations, through hands-on experiences, and so forth.

Endnotes

[1] Viadero, D. (1991). "Parents in S.C. Attack Alleged 'New Age' Program." *Education Week*.

[2] For reviews see: Reeves, L.M. and Weisberg, R.W. (1994), "The role of content and abstract information in analogical transfer," *Psychological Bulletin*, 115, 381-400; Barnett, S.M. and Ceci, S.J. (2002), "When and where do we apply what we learn? A taxonomy for far transfer," *Psychological Bulletin, 128* (4), 612–637.

[3] Chen, Z., Mo, L., and Honomichl, R. (2004). "Having the memory of an elephant: Long-term retrieval and the use of analogues in problem solving," *Journal of Experimental Psychology: General, 133*, 415–433.

[4] For a readable review see: Baron, J. (2000). *Thinking and Deciding*, Cambridge, UK: Cambridge University Press.

[5] Rotherham, A. J. and Willingham, D. T. (2010). "'21st Century' Skills: Not New, But a Worthy Challenge," *American Educator, 34* (1), 17–20.

[6] National Research Council (2007). *Taking Science to School,* Washington, D.C.: National Academies Press.

[7] Friedler, Y., Nachmias, R., and Linn, M. C. (1990). "Learning scientific reasoning skills in microcomputer-based laboratories," *Journal of Research in Science Teaching, 27,* 173-191.

[8] Schauble, L., Glaser, R., Raghavan, K., and Reiner, M. (1991). "Causal models and experimentation strategies in scientific reasoning," *The Journal of Learning Sciences, 1,* 201-238.

[9] For example see: Dunbar, K. N. and Fugelsang, J. A. (2005), "Causal thinking in science: How scientists and students interpret the unexpected," in M. E. Gorman, R. D. Tweney, D. C. Gooding, and A. P. Kincannon (eds.) *Scientific and Technological Thinking,* 57-79, Mahwah, N.J.: Erlbaum; Echevarria, M. (2003), "Anomalies as a catalyst for middle school students' knowledge construction and scientific reasoning during science inquiry," *Journal of Educational Psychology, 95,* 357-374.

[10] Taconis, R., Ferguson-Hessler, M.G.M., and Broekkamp, H., (2001). "Teaching science problem solving: An overview of experimental work," *Journal of Research in Science Teaching, 38*(4), 442-468.

This paper is adapted from "Critical Thinking: Why Is it So Hard to Teach?," which Willingham wrote for the summer 2007 issue of American Educator. *The sidebar is adapted from "Knowledge in the Classroom," which Willingham wrote for the spring 2006 issue of* American Educator.

Campaign Steering Committee

Leslye Arsht, Chief Executive Officer and Co-founder of StandardsWork;

Michael Casserly, Executive Director of the Council of the Great City Schools;

Ron Fairchild, Chief Executive Officer of the Smarter Learning Group and Senior Consultant for the Campaign for Grade-Level Reading;

Chester E. Finn, Jr., Distinguished Senior Fellow and President Emeritus of the Thomas B. Fordham Institute;

Kati Haycock, President of The Education Trust;

E. D. Hirsch, Jr., Founder of the Core Knowledge Foundation;

Joel Klein, Chief Policy and Strategy Officer of Oscar Health Insurance and former Chancellor of the New York City Department of Education;

Robert Pondiscio, Senior Fellow and Vice President for External Affairs of the Thomas B. Fordham Institute;

Ralph Smith, Senior Vice President of the Annie E. Casey Foundation and Managing Director for the Campaign for Grade-Level Reading; and

David Steiner, Executive Director of the Johns Hopkins Institute for Education Policy and former New York State Commissioner of Education.

Campaign Scientific Advisory Board

Marilyn Jager Adams, Visiting Scholar, Brown University;

Donald Compton, Professor of Psychology and Education, and Associate Director, Florida Center for Reading Research, Florida State University;

Anne E. Cunningham, Professor and Head Graduate Adviser, Cognition and Development, Graduate School of Education, University of California, Berkeley;

Nell K. Duke, Professor of Literacy, Language, and Culture, Combined Program in Education and Psychology, University of Michigan;

Claude Goldenberg, Nomellini & Olivier Professor of Education, Graduate School of Education, Stanford University;

David Klahr, Walter van Dyke Bingham Professor of Cognitive Development and Education Sciences, Department of Psychology, Carnegie Mellon University;

Danielle McNamara, Professor, Department of Psychology, Arizona State University;

Susan Neuman, Professor, Early Childhood and Literacy Development, Steinhardt School of Culture, Education, and Human Development, New York University;

Timothy Shanahan, Distinguished Professor Emeritus of Urban Education at the University of Illinois at Chicago; and

Daniel Willingham, Professor of Psychology, University of Virginia.

Knowledge Matters is a campaign to make building knowledge Job One for American education.

It's time to restore history, science, geography, art, and music to the education we give to all students, especially those least likely to gain such knowledge outside school. Greater comprehension, critical thinking, curiosity, and equality will be our reward.

Robert Pondiscio
Executive Director
Robert@KnowledgeMattersCampaign.org

Lisa Hansel
Director
Lisa@KnowledgeMattersCampaign.org