CHAPTER

70 Elements of Effective Learning

Jeffrey D. Karpicke *and* Garrett M. O'Day

Abstract
This chapter seeks to understand effective learning strategies at the level of principles of learning and memory, fundamental truths about how learning and memory operate. Learning and memory are adaptive abilities for coordinating actions in a complex environment. Viewing learning as coordination, rather than as the storage of knowledge and experiences, is essential for understanding effective strategies. Three *elements of effective learning* are described. Effective learning occurs when retrieval cues are available that permit people to express their knowledge (cue availability), when those retrieval cues are diagnostic of target knowledge (cue diagnosticity), and when elaborative study methods have prepared learners to use potential retrieval cues by promoting organization and distinctiveness (elaboration). Each element of effective learning is illustrated with examples from foundational research. The final section evaluates several strategies that leverage elements of effective learning and offers examples of how students might use the strategies in educational settings.

Key Words: learning, strategies, retrieval practice, spacing, interleaving, elaboration

Introduction

The charge for this chapter was to write about "principles of memory in the classroom." Cognitive science has provided a foundation of evidence about what works (and what does not) to improve learning and memory under a variety of conditions, for a variety of topics and materials, and across a variety of learners. During the past few decades, there has been a sharp increase in interest in applying cognitive science to education, and the evidence base continues to grow and expand. One challenge is condensing recent research and making it meaningful to educators, who would rightfully scratch their heads about how, for example, a single experiment on learning Swahili–English word pairs would have any relevance to their students learning about a variety of topics in their classrooms. What are the principles that govern how learning and memory work, and what are the implications of these principles for how students learn in school? Is it possible to connect these disparate dots?

One approach to discussing principles of memory in the classroom might be to line up and review the usual suspects—strategies that have received considerable attention in recent years. This chapter covers effective strategies such as retrieval practice, spacing, interleaving, questioning techniques, and others, and it also touches on a few ineffective strategies as well. The trouble is, these are not principles of learning and memory. Recommendations such as "interleave different types of practice problems" and "don't passively reread materials" are good advice, but these are statements about specific techniques and practices. *Principles* ought to describe universal and fundamental truths about the nature of how people learn and remember. Effective strategies should be grounded in principles of learning, but those strategies themselves are not principles.

Perhaps the trouble runs even deeper. Are there principles of learning and memory at all? Some scholars in educational circles take the position, essentially, that because there are individual differences among people, there are no principles of learning, or at least the search for principles would be a fool's errand. Coming from the cognitive psychology of learning and memory, Roediger (2008) made a similar argument. He proposed that the search for general laws of learning and memory, which guided research on learning during the middle of the 20th century, was abandoned largely because there are no laws of learning and memory. Perhaps there is a fine line between a "law" and a "principle." Regardless, Roediger's argument was that because all memory effects have exceptions—situations in which strategies such as retrieval practice, spacing, and interleaving do not work—there are not universal laws or principles of learning and memory.

This chapter proposes that there are principles of learning and memory, fundamental truths about how learning and memory operate (Surprenant & Neath, 2009), and that these principles have important implications for understanding and promoting student learning in school. Principles of learning may be hiding in plain sight because they are always at work, so there is value in stating them explicitly. This chapter identifies three guiding principles that are essential for understanding learning and memory. The field has known these principles for a long time: They were laid out more or less in Tulving's (1983) foundational book *Elements of Episodic Memory*. The principles of learning described here could be called the *elements of effective learning* because each element is essential for understanding and promoting learning. Yes, there are indeed situations in which some learning strategies do not produce their desired effects, but even in those cases, the core operating principles of learning and memory are at work.

Bridging the gap between foundational research on learning and memory and student learning in educational settings is a lofty goal, and this chapter is a small part of that larger project. The chapter describes some of the foundational work that identified principles of learning and memory in laboratory research with simple tasks such as remembering lists of words. The core principles of learning are also crucial for understanding complex learning situations that require learners to apply their knowledge and solve problems. Principles of learning can explain why certain strategies are effective and why learning is successful under some conditions but unsuccessful under others. Throughout the chapter, a few examples of students learning in school are used to illustrate principles and strategies in action.

Following this introduction, the second section of the chapter proposes that to understand some of the core principles, it helps to shift one's mindset about learning and memory. People have an everyday view or metaphor of what memory is and how it works. Most everyday metaphors characterize learning and memory as a system for recording and storing knowledge and experiences. This chapter argues that it is more accurate and useful to view learning and memory as adaptive abilities for *coordinating* actions in a complex environment. Rather than emphasizing recording or storage processes, this view of learning as coordination places the emphasis squarely on retrieval processes, the processes involved in drawing upon the past to meet the demands of a present situation.

The third section lays out three principles of learning and memory. The first principle is *cue availability*. All knowledge is expressed in a retrieval context, so the cues that are available and how learners interpret those cues are critical for what knowledge they express. The second principle is *cue diagnosticity*. Effective learning and memory performance happens when retrieval cues match desired target knowledge without matching other irrelevant knowledge. With these two principles in mind, the next question is, What activities would prepare learners to retrieve and apply their knowledge in a future retrieval context? The third principle is *elaboration*, which answers the question by stating that effective encoding activities add details that make knowledge distinctive and organized.

The fourth section of the chapter turns to learning strategies, evaluating several strategies that have received considerable attention in recent research. Effective strategies may be challenging and effortful, and sometimes the benefits of effective strategies are not evident during learning but become prominent in the long term. This pattern is known as creating "desirable difficulties" that enhance learning (Bjork, 1999), and it is a challenge because some effective strategies may improve long-term learning even though the benefits are not readily apparent in the short term (see Chapter 71). The final section of the chapter discusses how effective strategies leverage the core operating principles of learning and offers examples of how students might use the strategies in educational settings.

Learning as Coordination

The nature of learning can seem abstract and challenging to grapple with. That is why for thousands of years people have used metaphors to understand and reason about learning and memory. Many metaphors throughout history have emphasized how memories are formed—the initial *encoding* of knowledge—and how they are *stored* in one's mind. Most metaphors have been linked to recording technologies that existed at the time. For example, Plato and Aristotle described memory as a wax tablet: Knowledge might be inscribed in one's mind just as it would be inscribed on a wax tablet, a common writing device in ancient times, with firmer impressions leaving longer-lasting inscriptions. Memory has been likened to a library filled with books, a house with rooms in which items are stored, and a tape recorder or video camera (Roediger, 1980). In education, a physical building metaphor is often used to talk about the mind. Knowledge is *constructed* by learners, and instructional techniques can provide *scaffolding* to aid learners. Another metaphor that still lingers today is that memory operates like a digital computer, a device for recording and storing information. Indeed, this metaphor gave the field the terms encoding, storage, and retrieval.

Talking about metaphors of learning and memory is not a frivolous exercise in philosophy. Metaphors are important because they have the power to guide the way people think about the world (Lakoff & Johnson, 1980). The perspective or mindset a person takes about learning and memory colors how the person thinks it works and what strategies the person thinks would be effective. The everyday metaphor most people adopt can be referred to generally as a *storehouse* metaphor. People tend to think of memory as a place where knowledge is stored, and learning involves the recording of new knowledge in that storage system. Not much consideration is given to how knowledge is recovered when it is needed. The emphasis is on getting things "in memory" with the hopes of getting them out later.

Researchers, teachers, and students would benefit from adopting a different mindset about learning and memory, one that does not characterize learning and memory as a storehouse for recording knowledge. One reason for a shift in mindset is that there is ample evidence that people do not record and store copies of past events and knowledge. Instead, people use retrieval cues to reconstruct their knowledge about events and the world. A striking example of the reconstructive nature of memory comes from a decade-long study of people's memories for the terrorist attack in the United States on September 11, 2001 (9/11), a deeply significant and emotional event that one would assume would leave an indelible record in people's minds. Hirst et al. (2015) asked people to report on their memories of 9/11—the circumstances in which they learned about the event and details about the attack itself. People were surveyed shortly after 9/11 (roughly 1 week to 10 days after the event) and completed follow-up surveys 11 months, 25 months, and 119 months (roughly 10 years) after the event. Remarkably, many details people reported on the original survey were forgotten during the first year. Even for this emotionally charged event, people did not record and store a mental copy of the event. Instead, people reconstructed their knowledge each time they attempted to remember, using the retrieval cues available to them at the time.

There are additional reasons to shift away from a storehouse mindset. Our learning and memory systems were not designed to record and store copies of past knowledge and events. The past never repeats itself—every situation, every retrieval context, is by definition new. A system that stores exact copies of the past would be of little help for adapting in a complex and dynamic environment. But perhaps the most important reason to abandon the storehouse mindset is because it places a premium on encoding and storage of knowledge and largely ignores retrieval processes. Shelving books in a library, filing documents in a file cabinet, placing items in a house, and building an edifice with the aid of scaffolding are all metaphors that emphasize encoding. A storehouse mindset has little to say about how knowledge is reconstructed and applied when it is needed in a particular retrieval context. And a storehouse mindset likely leads students to adopt ineffective learning strategies. If getting knowledge "in memory" is what matters, then students may adopt strategies such as repetitive reading because they think that increasing the sheer exposure to information will create a deeper impression, like repeatedly engraving onto a mental wax tablet.

The alternative mindset is to view learning and memory as an adaptive system designed to help people *coordinate* their actions in a dynamic environment (Schwartz & Goldstone, 2016). Learning is the ability to use the past to meet the demands of the present. People use the cues available to them in a current retrieval context to reconstruct and apply their knowledge. When people are asked to remember details of an event, such as where they were on 9/11, they use that cue along with other cues generated as they search memory to reconstruct what they think happened at that time (Kahana, 2017). Likewise, when a student is asked to apply their knowledge or solve a novel problem, the problem is a cue to draw upon prior knowledge that might lead to a solution. In both scenarios, people use retrieval cues to reconstruct their knowledge as a means of coordinating their actions and meeting the demands of the task at hand.

A coordination viewpoint fits with a variety of metaphors—examples of systems that can produce an outcome without recording and storing copies of it. A good example is a musical instrument such as a piano (Wechsler, 1963). Pianos are capable of producing melodies, but melodies are not "stored" inside pianos. A piano has the capability to produce melodies, given that it is tuned a certain way and given that keys are pressed in a particular pattern, but it would be strange to ask how or where melodies are stored inside a piano. The nervous system generally works in a similar way to allow people or other animals to coordinate their actions in the environment. The visual system provides the ability to see objects but does not need to store copies of objects to do so. Sensory systems allow people to experience sensations such as cold, pressure, or pain. When a person's arm is pinched, the person may experience pain, but it would be strange to think pain was stored somewhere in the arm. Likewise, learning and memory systems are tuned by experiences and give people the ability to use retrieval cues to reconstruct knowledge and thereby coordinate their actions in the environment.

When learning is viewed as coordination, rather than as the recording and storage of knowledge, the emphasis and crucial questions shift toward retrieval—specifically about the conditions that create effective coordination. The questions become, for example, What is the context in which learners will need to reconstruct and use their knowledge? What would make potential retrieval cues effective? And what could learners do to prepare for an anticipated retrieval situation? The answers to these questions provide a framework for the elements of effective learning.

Elements of Effective Learning

To ground this discussion of principles of learning, Box 70.1 presents three examples of students at different grade levels learning different topics in school by engaging in different learning activities. In each case, what are the available retrieval cues in the learner's environment? What would make those cues effective for reconstructing and applying knowledge? Are the students using effective learning strategies? And what learning strategies or instructional techniques would prepare these students for success? This section describes foundational research on three principles of learning and returns to these student cases to view them through the lens of each principle.

Principle 1: Cue Availability

The only way to examine what a person knows or can do is by having them engage in retrieval processes. There is no way to assess a person's knowledge without retrieval. Therefore, the first principle for understanding learning is referred to as cue availability. A student does not answer a question until that question is asked or solve a problem until the problem is posed, just as a piano does not play any melodies until the keys are pressed. A student may have the *potential* to recall facts, answer questions, and solve problems, but the only way to witness that potential is by asking the person to engage in some form of retrieval. Therefore,

Box 70.1 Three Illustrative Cases of Students Learning Academic Content in School

Mason is a fourth grader who is learning about seasons and different weather conditions in his current science unit. He reads short articles about weather phenomena and fills out worksheets. For example, one worksheet requires him to match key terms such as tsunami, thunderstorm, and drought with images (a huge wave, clouds and lightning bolts, and sun and desert). His teacher leads classroom discussions and asks the class questions.

Josephine, an eighth-grade student, is learning about geometry—topics such as calculating the volume and surface area of objects. Her teacher explains and demonstrates how to solve problems, and she takes notes in class. Her homework involves solving practice problems for which she can refer to her notes and worked examples from class.

Kim, a high school junior, is learning about World War II in her Advanced Placement United States History class. She attends class and takes notes, and she is reading a textbook chapter about the war in the Pacific to prepare for an upcoming exam.

it is essential to consider the retrieval cues available in a particular context and how students interpret those cues, because the knowledge a person expresses can vary greatly depending on retrieval conditions.

The foundational research that put retrieval cueing on the map was done by Tulving and Pearlstone (1966). In a very large laboratory experiment, they had students study lists of unrelated words such as *cow*, *bomb*, *radio*, and *pepper*. They examined several conditions, but it will suffice to describe just two. One group of students studied a list of 48 words and then freely recalled them, writing down as many as they could in any order. Another group studied the same list but then recalled them with the aid of retrieval cues, words that had not been on the list but that were aimed at reminding the subjects of the original words (such as *animal* as a cue to remember *cow*, *weapon* to remember *bomb*, *entertainment* to remember *radio*, and *food* to remember *pepper*). The average number of words recalled is shown in Figure 70.1A. Whereas subjects recalled roughly 19 of the 48 words on the free recall test, they could recall approximately 36 words on the cued recall test.

What's the big deal? It may not be terribly surprising that when you give people hints or clues, they recall more items than they do without hints. But Tulving viewed the results from a different perspective and asked an incisive question: What had happened to the 17 words that were recallable on the cued recall test but not in free recall? Why had subjects forgotten those words on the free recall test? The explanation could not be about encoding or storage, because the two groups had studied and presumably stored the words in the same way. The only way to explain the difference was by considering retrieval processes, and this led to Tulving's major insight. People have a vast amount of knowledge available in their minds and memories, or a vast potential to produce knowledge, but they only access a portion of it at any given moment. The factors that determine access are the retrieval cues available in a particular retrieval context. The subjects who freely recalled 19 words may very well have had 17 more "recorded" and "stored" in mind, but they could not express that knowledge without the availability of

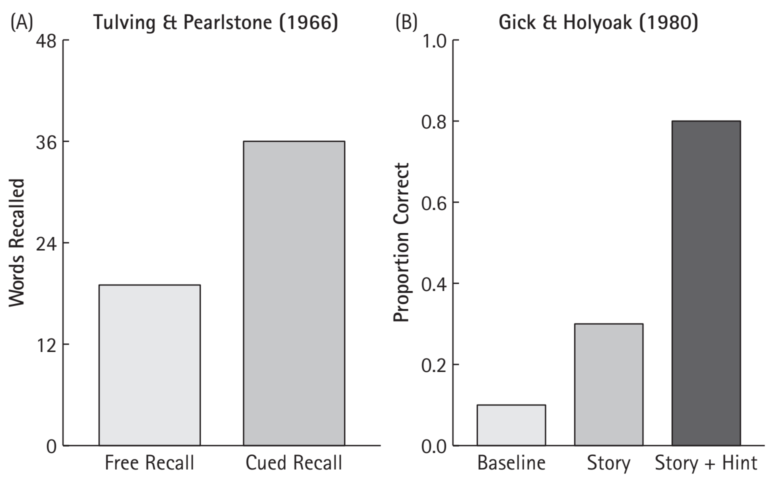

Figure 70.1 The importance of cue availability and interpretation. (A) Results from Tulving and Pearlstone (1966), showing that the number of words recalled increased under cued recall conditions relative to free recall. (B) Results from Gick and Holyoak (1980), showing that problem-solving performance increased dramatically when students were told to think back to a prior study (a hint) relative to when they were not.

additional cues. People have a vast potential to produce knowledge, but the portion expressed depends on the cues available in a particular retrieval context.

In the interest of connecting disparate research areas, let us turn now to a different learning scenario. Consider the following problem (paraphrased from the original):

> A doctor wants to use rays to destroy an inoperable tumor inside a patient's body, but also wants to prevent the rays from destroying healthy tissue. High-intensity rays are needed to destroy the tumor, but those would also destroy healthy tissue. Low-intensity rays would not destroy healthy tissue but also would not affect the tumor. How could the doctor destroy the tumor with rays but avoid destroying the healthy tissue?

In this famous problem, originally used in research by Duncker (1945), the solution is to divide the ray into multiple low-intensity rays that converge on the tumor, thereby destroying the tumor without destroying surrounding tissue. In a landmark paper on problem-solving, Gick and Holyoak (1980) found that 10% of students came up with the convergence solution to the radiation problem on their own without any additional information. Students in another condition read an analogous story before solving the problem. The story was about a general who, in order to conquer a fortress, had to divide his army into multiple small groups that converged from different directions on the fortress. The analogous story provided a narrative example of a convergence solution that could be helpful for solving the radiation problem. When students read the analogous story prior to solving the problem, approximately 30% produced the convergence solution. But a remarkable thing happened when students read the analogous story and were also told to think back to the original story when they solved the problem. Now, as depicted in Figure 70.1B, approximately 80% of students produced the convergence solution.

How can this dramatic improvement in problem-solving performance be explained? A considerable amount of research on problem-solving has focused on how learners encode the deep structure of materials, or schemas, that might be used to solve future problems (Gentner & Maravilla, 2018). Reading an analogous story improved problem-solving relative to not reading it because the story gave learners an example or schema of the convergence solution. But the Gick and Holyoak (1980) results, especially when juxtaposed with Tulving and Pearlstone's (1966) results, illustrate just how powerful the retrieval context is in problem-solving performance. Many students had indeed understood the deep structure of the analogous story and could apply it to solve the radiation problem. They possessed the necessary knowledge to solve the problem, and just as in the Tulving and Pearlstone experiment, the critical difference across conditions was not about the encoded knowledge or abstracted schema. The largest effect on problem-solving occurred when the retrieval conditions—and only the retrieval conditions—were changed.

Cue availability is a crucial element of learning because all inferences about what a person knows depend on the cues provided in a retrieval context. Knowledge cannot be examined in a vacuum, devoid of any retrieval context. Gick and Holyoak (1980) could not open their subjects' heads and look directly at an encoded schema. The only way to examine learning is by having people engage in some form of retrieval processes. When students' knowledge is assessed in school, they are given retrieval cues and asked to reconstruct their knowledge to coordinate their actions in a retrieval context, whether that means recalling facts, answering questions, writing essays, or solving novel problems. With that in mind, consider the retrieval contexts and cues that might be available to the students described in Box 70.1.

Kim, the high school student taking U.S. history, is expected to know about historical events, key people, and cause–effect relationships. How will her knowledge be assessed?

In many classrooms at many levels of education, knowledge is assessed on exams and quizzes. Thus, Kim's opportunity to express her knowledge in school will be mainly on classroom exams that include multiple-choice questions and short essays. Likewise, Josephine, the eighth-grade student learning about geometry, will have classroom exams on which she will be asked to solve problems, for example, by determining the surface area and volume of shapes such as circles and spheres. For Mason, the student in fourth grade, the stakes attached to classroom quizzes may not be as high as they are in middle or high school, but classroom quizzing begins in elementary school. His teacher might give him a quiz with questions about facts that the class has learned about different types of weather. Knowing about the conditions in which knowledge will need to be retrieved is the first step in knowing how to prepare.

Learners also reconstruct and use their knowledge in a wide range of situations outside the classroom. Kim might think about important dates or past historical events when she reads or hears about current events. Josephine might think about surface areas when she is arranging things on the walls of her bedroom or about volumes when she is packing a backpack or suitcase. Mason might think about things he learned in school when he hears a weather report on the news in the morning or when he notices different types of clouds in the sky on different days. In each situation, the ability to reconstruct and use knowledge depends on the retrieval cues available in a specific context.

No single assessment is a perfect or complete indicator of a student's knowledge. A person might not be able to express knowledge in one situation but might be successful in another, just as students in Gick and Holyoak's (1980) experiment did not always solve the radiation problem without a hint but were capable of solving it when prompted with a hint. Retrieval processes are variable, as demonstrated in the variability of each person's repeated reports about the details of 9/11 (Hirst et al., 2015). The same phenomena exist in classroom settings. Many people have had the frustrating experience of not knowing the answer to a question during an exam and having the answer come to mind later, in an entirely different context. Teachers dislike standardized tests for many reasons, but a common complaint is that the tests do not truly assess what students know. In other words, students would be able to express their knowledge if they were assessed a different way, with a different retrieval context and cues. All of these facts illustrate just how critical retrieval processes are for understanding learning.

The particular cues available in a learner's environment matter crucially for what they will remember and what knowledge they will express. A student may possess some ability, skill, or knowledge but not express it in a particular retrieval context. This leads naturally to the next principle. When cues are available to a learner, what is it that makes those cues effective?

Principle 2: Cue Diagnosticity

The idea of cue diagnosticity can be illustrated with an example. Imagine we were thinking of a person and giving you clues to help you identify the person. We could tell you that the person is a man who is tall, intelligent, wears a suit to work, and is a good speaker in front of audiences. Those are 5 features of this person but probably not enough to guess the person's identity. We could add that the person is famous (although you probably inferred that), is married and has children, is often considered to be charismatic and a leader, and is known to play basketball. Even with roughly 10 features now, you still may not be able to guess the person's identity. But had we not told you any of those features and instead just provided a single feature—that the person was the 44th President of the United States—you may have been very likely to guess that the person was Barack Obama. In the first part of this example, the retrieval cues contained several features that matched the target person. But those cues and features also matched many other potential people, so they may have been unlikely to lead you to think of him. In the second example, the retrieval cue was a single feature that uniquely matched Barack Obama. Even though that cue provided less of a total match, it was probably more effective.

A retrieval cue is effective when it is *diagnostic* of target information—when it contains features that match a target but that do not match many other possible targets. The availability of retrieval cues and how people interpret those cues are essential elements of learning. But even when retrieval cues are available, performance can be quite variable. Learners will be in the best position to reconstruct and apply their knowledge when they have diagnostic cues that uniquely specify target knowledge. The crucial question then becomes what makes retrieval cues effective. Why would one set of retrieval cues lead to poor performance when a different set of cues would create a better outcome?

Tulving carried out a program of research that established foundational knowledge about what makes retrieval cues effective. In one experiment, Thomson and Tulving (1970) had people study a list of target words, such as *chair*, shown individually or paired with another word (*glue–chair*). Later, the subjects were asked to recall the words. In one retrieval condition, the subjects were given cues that were obviously related to the targets (e.g., *table* as a cue to recall *chair*). As shown in the left portion of Figure 70.2A, these cues with strong meaningful relationships to the targets, known as strong associates, were effective for recalling individual words. However, the situation changed dramatically when instead of studying individual words, the subjects had studied the targets paired with weak associates, words such as *glue* that did not bear strong semantic relationships to targets such as *chair*. Now, as shown in Figure 70.2A, *glue* was an extremely effective cue to remember the word *chair*, even though the words share little meaningful relationship. The other striking feature of the results was that the strong associates (*table*) that were so effective for recalling individual words (*chair*) were extremely ineffective when the targets had been paired with a different word during the study phase (*glue–chair*).

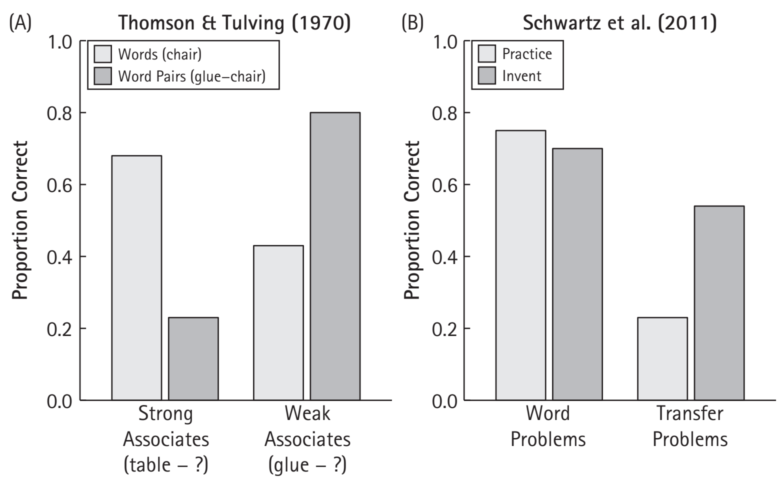

Figure 70.2 The importance of cue diagnosticity. (A) Results from Thomson and Tulving (1970), showing that the effectiveness of studying words or word pairs depended on how retention was assessed—that is, the diagnostic value of retrieval cues. (B) Results from Schwartz et al. (2011), showing that the effectiveness of two different instructional methods depended on how problem-solving performance was assessed.

The results reported by Thomson and Tulving (1970) and others like it were groundbreaking because they showed that mere similarity or semantic relatedness of cues and targets were not the key factors that make cues effective. Instead, the important factor for making cues effective was when those cues help people reinstate the initial learning context. Although *glue* does not share a strong relationship to *chair*, when that word was part of the original study experience, it became an extremely effective retrieval cue later on. Tulving referred to this as the *encoding specificity principle* (Tulving & Thomson, 1973), which states that part of what makes a cue effective is that it matches and thereby helps a person reinstate a study experience.

It is important to emphasize, however, that this represents only part of the explanation for the effectiveness of retrieval cues. In experiments such as Thomson and Tulving's (1970), the cue *glue* "matched" the target word *chair* but also did not match any other words on the list. If *glue* had been paired along with all of the to-be-learned words, it would match many possible targets and would lose its effectiveness, just as cues such as "tall, intelligent, and wears a suit to work" were poor cues for Barack Obama because those cues match many other possible people. This situation is referred to as cue overload: Retrieval cues lose their potency when they match many possible targets (Nairne, 2002, 2006; Watkins & Watkins, 1975). Therefore, retrieval cues become diagnostic and effective when they uniquely match target knowledge without matching several other possibilities (for related discussion, see Chapter 40).

The Thomson and Tulving (1970) results illustrate factors that make retrieval cues effective in a laboratory setting, but it is crucial to emphasize that cue diagnosticity is operating in all learning situations, and it is readily apparent in a great deal of educational research. One illustrative example comes from Schwartz et al. (2011). They worked with eighth-grade students who were learning about physics topics such as density, speed, surface pressure, and the spring constant. The deep structure common to all of these topics is the *ratio* between physical properties (density is mass over volume, speed is distance over time, surface pressure is force over area, and the spring constant is force over displacement). The procedure used by Schwartz et al. was complex and spanned multiple days of instruction, so this summary emphasizes key differences between conditions. In a *practice* condition, students were told how to calculate density and then practiced solving problems with a combination of worked examples (step-by-step examples of how to solve a problem) and cases (example problems for students to solve). This condition was aimed at reflecting common instructional practices in many classrooms. In a second condition, called an *invent* condition, students were given a similar set of examples and cases to work with and told to "invent an index"—that is, come up with a measure that captured the relationship shown in the examples and cases. The researchers assessed knowledge in several ways, but two methods make the key point for present purposes. On a word problem test, students were given word problems that required them to use their knowledge about how to calculate density. On a second test, referred to here as a transfer test, students were given problems that asked them to determine the stiffness of the fabric of different trampolines with varying numbers of people standing on them. The trampoline problems tapped into knowledge about the spring constant, and to answer those problems, students needed to draw upon the deep-structure knowledge that fabric stiffness would be a ratio (number of people over displacement).

Students' performances on the word problems and the transfer test (the trampoline problems) are shown in Figure 70.2B. On the final word problem test, the two instructional conditions performed roughly the same. In fact, the practice condition performed slightly better than the invent condition (approximately 5%), although the difference did not reach significance. Based on these data alone, one might conclude that the two instructional conditions were equally effective or perhaps that the practice condition may have produced slightly more learning. However, on the trampoline problems, shown in the right portion of the figure, the invent condition produced substantially better performance relative to the practice condition.

These two disparate examples—recalling words in a laboratory experiment and solving physics problems in an eighth-grade classroom—both point to the essential importance of cue diagnosticity. The available cues and the diagnosticity of those cues were crucial for the knowledge that learners expressed. In the Schwartz et al. (2011) study, the word problems asked students to use their knowledge about solving density problems, and both instructional conditions had prepared students equally well to apply their knowledge in that context. But the trampoline problems required students to access deep-structure knowledge about ratios, and the invent-an-index instructional condition prepared students more effectively for that retrieval scenario. The pattern in both the Thomson and Tulving (1970) and Schwartz et al. (2011) experiments is known as an *encoding–retrieval interaction* (Roediger et al., 2017; Tulving, 1983). Which of the two encoding conditions or instructional methods was more effective? The answer is that it depends on the retrieval situation, specifically on the cues that are available and the diagnostic value of those cues.

For each of the student learning cases in Box 70.1, how diagnostic are the retrieval cues the students have at their disposal when their knowledge is being assessed? Retrieval cues and contexts may match a variety of knowledge. The best performance will happen when the available cues uniquely match relevant knowledge that the learner needs to answer questions or solve problems.

In Kim's high school U.S. history course, portions of her classroom exams contain multiple-choice questions such as the following: "Which of the following battles marks the turning point of World War II in the Pacific Theater? A) Pearl Harbor, B) Iwo Jima, C) Okinawa, or D) Midway." All of the possible answers are plausible and refer to locations where important battles took place. Thus, "World War II battles" is an overloaded cue—it could specify multiple possible answers. The question also asks about the Pacific Theater and "turning point," additional features that make the cue more diagnostic of specific target knowledge. For Kim, the question will be diagnostic of the correct answer (Midway) if she has knowledge about what happened in each of the battles listed in the question, when the battles occurred, and how important the battles were in the course of the war. As discussed in the next section, a knowledge scheme that is organized and contains distinctive information would help Kim answer this question, essentially turning the question into a diagnostic retrieval cue.

In Josephine's geometry class, she learns how to calculate the surface area and volume of several different shapes. Thus, the potential retrieval cues *surface area* and *volume* are linked to multiple formulas for a variety of objects. A question about calculating the area of a circle is a cue to retrieve and apply a particular formula, but Josephine needs to be able to discriminate which particular formula would solve the particular problem. Cue diagnosticity is critical because thinking of the formula to calculate the surface area of a sphere instead of the surface area of a circle would result in an incorrect answer. From the standpoint of cue diagnosticity, Josephine will be able to solve problems successfully when those problems (the retrieval cues) bring to mind the necessary formulas without bringing to mind other possible formulas.

For Mason, consider a few questions he might have on a worksheet or quiz, such as the following: "What type of cloud is thin, wispy, and forms high above the ground?" Perhaps the importance of cue diagnosticity is obvious by now. This question would serve as a diagnostic retrieval cue for Mason when the features *thin*, *wispy*, and *forms high above the ground* specify a specific type of cloud (a cirrus cloud) and not another type (e.g., stratus or cumulus clouds). Another question might be, "What types of dark clouds are likely to be in the air during a thunderstorm?" Again, the question is an effective cue when it brings to mind certain types of clouds (cumulus or cumulonimbus) but not others (cirrus or stratus).

It is crucial to anticipate the retrieval environment—the situation and conditions in which a learner will need to use and apply their knowledge. Learning strategies will be effective when they prepare the learner for that situation. Two elements are in now place: Cues need to be available, and effective cues are diagnostic and uniquely specify target knowledge. The last of the three elements asks, What are the features of learning and instructional strategies that prepare learners for success in a given retrieval context?

Principle 3: Elaboration

"To elaborate" means to add detail, and in basic memory research, elaboration refers to encoding activities that require a learner to add detail to a knowledge representation. Ideas about elaborative encoding were introduced in foundational papers on depth of processing approximately 50 years ago (Craik & Lockhart, 1972; Craik & Tulving, 1975; Craik & Watkins, 1973; Tulving & Madigan, 1970). In a retrospective review, Craik (2002) explained elaboration and why it would benefit learning and memory:

> Why should greater trace elaboration support good retention? Two possibilities are, first that a richly elaborate trace will be more differentiated from other episodic records—this greater distinctiveness in turn will support more effective recollection in an analogous way to distinctive objects being more discriminable in the visual field. A second (complementary) possibility is that elaborate traces are more integrated with organised knowledge structures which, in turn, serve as effective frameworks for reconstructive retrieval processes. (pp. 306–307)

Craik's description highlights two critical aspects of elaboration. An elaborative activity is thought to be effective because it makes knowledge more *distinctive*, but elaboration also helps learners create *organized* knowledge structures. Importantly, both of these factors—organization and distinctiveness—are essential to prepare learners for future retrieval situations in which they will need to reconstruct, use, and apply their knowledge (Hunt & McDaniel, 1993).

An example of the importance of both organization and distinctiveness in elaboration is provided by Hunt and Smith (1996). They had subjects study a list of 50 words in which groups of 5 words were related (e.g., *herring*, *bluegill*, *trout*, *guppy*, and *catfish*). In a similarity-encoding condition, subjects were told to write down one thing that was common among all 5 words in a group (here, the words are all types of fish). In a difference-encoding condition, subjects were told for each word to write down one thing that might be related to that word but not related to others (in this example, a subject might write "pickled" as a unique word for *herring*). The left portion of Figure 70.3A shows that when subjects were asked to freely recall the words, having generated distinctive aspects of each word produced better recall relative to generating shared aspects for all the words. The benefit also occurred in a condition in which subjects were given back the words they had produced and could use them as retrieval cues. Note that in this condition, subjects were able to recall nearly the entire list perfectly (average recall was 97%; see also Mäntylä, 1986).

It is important to emphasize that effective elaboration will involve *both* organization and distinctiveness. In the Hunt and Smith (1996) experiment, the similarity of the sets of words was obvious to learners, and because the organization was obvious, the distinctive encoding task was especially effective. In cases in which the organizational structure of materials is less obvious (e.g., imagine a list of completely unrelated words), elaboration strategies that emphasize similarity are important and effective. The materials a person is learning may afford

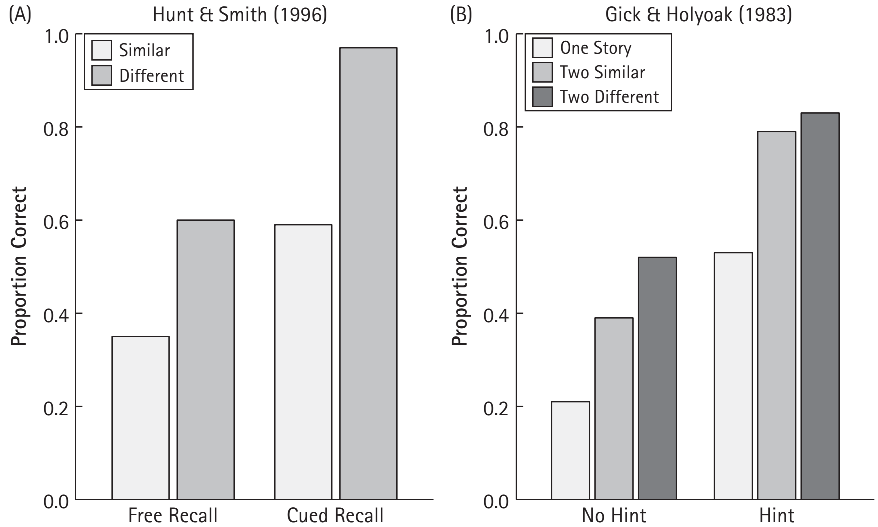

Figure 70.3 The importance of elaboration. (A) Results from Hunt and Smith (1996), showing that elaborating by generating distinctive cues while studying enhanced later retention. (B) Results from Gick and Holyoak (1983), showing that experiencing multiple distinct analogous stories, which afforded elaborative encoding, tended to enhance later problem-solving performance.

organizational or distinctive encoding, and the best learning occurs when elaborative strategies complement the materials so that learners encode knowledge that is both organized and distinctive (McDaniel & Einstein, 1989, 2005).

Once again, principles of learning are relevant well beyond the realm of simple word list experiments, and the influential research on problem-solving by Gick and Holyoak (1983) provides another excellent example. In one of their experiments, subjects attempted to solve the radiation problem (described on p. 15) after reading one analogous story or two analogues that were similar or different. Students in the similar-encoding condition read two analogues either about a military campaign (like the story about the general described earlier) or about firefighting, in which firefighters use multiple converging hoses to put out a fire. Students in the difference-encoding condition read one analogue of each type. The problem-solving results are shown in the right panel of Figure 70.3B. Students benefitted when they studied two story analogues relative to studying only one, and the results also depended on the similarity of the two stories and on the retrieval conditions—whether students solved the problem with or without a hint.

The core principles of learning—principles about elaboration and retrieval cuing—can provide an explanation for this pattern of results. Studying multiple stories improved the encoding of the deep structure common across the stories (the convergence solution). Studying different distinctive stories improved students' ability to interpret and recognize the radiation problem as a retrieval cue to recover relevant knowledge when they were not given a hint. But when given a hint, both groups that had studied two analogous stories exhibited the same level of problem-solving performance. Those two groups had encoded the deep structure of the materials equally well, but students in the difference-encoding condition were better prepared to use that knowledge when the retrieval context did not include an explicit prompt to think back to their prior experiences.

Returning to the examples of student learning in Box 70.1: Are the students in each scenario engaging in elaborative study methods that might help them reconstruct and use their knowledge in the future when they need it?

The conditions in which Kim would need to retrieve and use her knowledge (her classroom exams) were described earlier. To study effectively, she should engage in elaborative strategies that help her organize knowledge (e.g., the order of historical events and cause–effect relationships) and also make that knowledge distinctive (unique features of historical events and people that differentiate them). Kim takes notes in class and completes assigned readings, which are markers of a conscientious student, but neither of these activities is especially elaborative. Specifically, neither activity by itself will add enriching detail to the material Kim is learning. The evidence-based strategies described in the next section would help Kim improve her preparation for her classroom exams.

Likewise, Josephine needs to know how to solve problems in her geometry class. Her classroom and homework activities involve solving practice problems, which is good preparation because it gives her practice performing the exact problem-solving skills that will be assessed on classroom quizzes. But there are a few relatively simple ways to enhance these activities that could be implemented by her teacher or by Josephine herself. Specifically, Josephine's preparation can be enhanced by improving *how* she goes about solving problems, *when* she solves the problems, and how the practice problems are *arranged* in classroom and homework assignments.

Mason, the student learning about weather phenomena, is in an active fourth-grade classroom. His teacher presents new content, asks questions, and leads classroom discussion, and Mason has many opportunities to think and talk about ideas with his teacher and with other students. As noted earlier, his knowledge about weather phenomena will be assessed on a quiz, and he has brief readings and worksheets with questions to complete as homework. Once again, there are simple and effective techniques that could be done both in the classroom and on homework assignments that would enhance Mason's learning.

Summary

The three principles described in this section provide an overall framework for looking at learning by asking what the ultimate retrieval situation will look like (cue availability), what conditions would lead to successful retrieval in that situation (cue diagnosticity), and what students can do in advance to prepare themselves for a retrieval situation (elaboration). Learning strategies can be evaluated within this framework—the elements of effective learning—by assessing how well specific strategies help learners prepare for successful performance in a future retrieval context.

Evaluating Learning Strategies

A wealth of research has evaluated the effectiveness of several learning strategies and instructional techniques that could be used in a variety of classroom settings. A full review is beyond the scope of this chapter (for additional discussion, see Chapter 71). Indeed, an entire practice guide commissioned by the Institute of Education Sciences covered 7 strategies (Pashler et al., 2007), an entire book-length review by Fiorella and Mayer (2015) covered 8 strategies, and an extensive review article by Dunlosky et al. (2013) covered 10 strategies. But even those publications only scratched the surface. John Hattie maintains a website, https://visible-learning.org, that is a meta-analysis of meta-analyses (you read that correctly) containing effect size estimates for more than 40 learning and instructional strategies and a total of 252 factors that influence student achievement.

This section examines six learning strategies that have received recent attention in cognitive and educational research. The first, repetitive reading, is an ineffective strategy that nevertheless remains popular among students. The other five are effective strategies founded on solid bases of research. For each strategy, a brief research example is provided, the strategy is evaluated in light of the elements of effective learning, and examples of how students and instructors might use the strategies are provided.

A theme that crops up across a number of strategies is delayed gratification: In some cases, the long-term rewards of engaging in effective strategies may not be immediately evident. Bjork (1999) has referred to this as creating *desirable difficulties* that enhance learning. Some strategies may make progress during initial learning slower and more effortful, but these strategies produce durable learning that lasts over the long term and transfers to novel contexts. Three strategies described in this section (retrieval practice, spacing, and interleaving) can at times create this particular pattern of results. Strategies that create desirable difficulties present a conundrum for students and educators because even though the strategies bolster long-term learning, they may not appear effective right away in the short term.

Repetitive Reading

Repetitive reading is one negative consequence of adopting a "storehouse" mindset about learning and memory. If a learner's goal is to get new knowledge "in memory," then repeated exposure to the material might increase the likelihood that knowledge is recorded and stored. Repetitive reading, however, is not an effective learning strategy, yet it remains extremely popular among students. In one survey that asked college students to list strategies they used when they studied, 84% of students said they repeatedly read their notes or textbooks (Karpicke et al., 2009). Repeated reading was almost twice as prevalent as the next most frequent strategy (solving practice problems, which was listed by 43% of students). More recent surveys indicate that students still report repeated reading of their notes and textbook as the most prevalent study habits (Blasiman et al., 2017).

If repetitive reading did indeed improve learning, then students would be making an effective choice by using the strategy. Unfortunately, that is not the case. Dozens of studies have shown that rereading materials can produce little or no benefit for students' retention of facts, let alone their deeper comprehension or ability to apply their knowledge and solve problems (Dunlosky et al., 2013). In a rigorous examination of the effects of rereading on learning, Callender and McDaniel (2009) had college students read lengthy educational texts such as those they encountered in their courses. The students took assessments that involved writing summaries, answering short-answer and multiple-choice questions, and that occurred immediately or after a delay. The manipulation in the experiment was simple: Students read the material once or twice. Across different texts, different assessment methods, and different retention intervals, there was no benefit of immediate repetitive reading on learning.

There is some evidence that rereading can be effective when repeated readings are spaced, for example, with a week delay between reading sessions (Rawson & Kintsch, 2005). Spaced learning is effective (as described below), but there are many more effective strategies that students could use during spaced study sessions other than rereading. Essentially, passive repetitive rereading lacks the type of elaborative processing necessary to improve learning and memory. When students simply reread material, there are likely little or no cognitive processes happening that help students go beyond the material in front of them, for instance, by adding distinctive details to ideas and concepts or by creating an organizational framework that would support long-term learning.

If Kim simply rereads her notes or her textbook when she studies for her U.S. history class, her mental organization of the material is unlikely to change, and she is unlikely to encode any unique or distinctive attributes about the historical figures and events she needs to know. Repeated reading is a recipe for poor performance. When it comes time to answer questions on an exam, Kim is likely to struggle if she only repeatedly read while she studied. She would have done little to improve the likelihood that the available retrieval cues will be diagnostic of target knowledge. Yet a beguiling aspect of repetitive reading is that it may increase students' fluency of processing as they reread, leading to the feeling that learning is happening. For all of the reasons stated here, it is best to dump rereading as a study strategy. A variety of other strategies are straightforward to use and far more effective than repetitive reading.

Teaching and Explaining

One effective strategy is learning by *teaching and explaining*, which involves at least two distinct stages that promote learning: preparing to teach and the act of teaching itself. When learners prepare to teach, they must first identify key pieces of information, organize that information in a coherent manner, and generate explanations of the material. When learners actually engage in the process of teaching the material to a peer, they must retrieve much of what they have prepared and generate explanations for important concepts in response to any questions they are asked. Thus, learning by teaching and explaining leverages multiple component processes that enhance learning.

Fiorella and Mayer (2014) attempted to distinguish the benefits of preparing to teach from the benefits of actually teaching. In their experiment, undergraduate students studied a short physics lesson about the Doppler effect. Students were told in advance that they were learning the lesson either in preparation for an upcoming test or in preparation to create a video-recorded lecture (i.e., preparing to teach vs. not preparing to teach). Within each of those conditions, half of the students only studied the lesson for 15 min, and the other half studied for 10 min and then created a video-recorded lecture in which they taught the material for 5 min (i.e., actually teaching vs. not actually teaching). Thus, the experiment included four conditions that disentangled the process of preparing to teach from the act of teaching itself. All students took a final comprehension test 1 week later that targeted students' ability to explain key concepts from the lesson. Overall, students who taught the material outperformed students who had not engaged in any teaching or explaining by creating the video-recorded lecture. Critically, the best performance was achieved by students who both prepared to teach and then actually taught.

The effectiveness of teaching can be explained in terms of the elements of effective learning. The first element of learning described earlier was cue availability. When students prepare to teach, they must anticipate the future retrieval context—the situation in which they will need to teach and explain—for example, by anticipating the order in which they will talk about ideas and anticipating what questions they might be asked about the material. Preparing to teach also promotes elaboration because students must develop an organized, coherent structure of the information in order to teach it to someone else. Students also must elaborate on individual ideas and concepts by creating explanations, which promotes distinctive encoding. Finally, when students engage in the act of teaching, they must further refine their mental organization of the material, generate novel explanations during interactions with their peers, and practice the act of retrieval itself.

Kim could prepare for her history exam by teaching and explaining a section of her history textbook to one of her friends in her class. To prepare to teach, Kim would need to do more than simply read the material; she would need to identify key ideas, people, and historical developments and organize this information into a coherent scheme that would aid her in explaining the material. For instance, she might group historical events in a cause-and-effect order, by geographical region, or by topic (e.g., war, economy, and politics). Explicitly organizing material while preparing to teach promotes elaborative processing. Kim could elaborate on the material further by creating explanations that she would use when teaching the content—explanations that would help make individual ideas within the material more distinctive. Finally, when Kim goes through the process of teaching and explaining material to a friend, she retrieves and reconstructs her knowledge as she is teaching. The act of retrieval is itself a powerful learning activity, as described in the next section.

Retrieval Practice

Retrieval practice is an effective strategy that can improve student learning for a wide range of educational content. Retrieval practice has received considerable attention in recent cognitive and educational research. Recent work has generalized the benefits of retrieval practice across a range of student populations, across a variety of materials and topics, and across assessments that require students to make inferences and apply their knowledge to novel problems (for a brief review, see Nunes & Karpicke, 2015; for a more comprehensive review, see Karpicke, 2017). Teachers can incorporate retrieval opportunities in their classrooms in a variety of ways, and it is fairly easy for students to practice retrieval when they study on their own, although students do not always engage in retrieval in the most effective ways.

One example of the effectiveness of retrieval practice comes from Roediger and Karpicke (2006). They had college students read brief educational texts on science topics and examined a few different study conditions, two of which will suffice for present purposes. In a repeated reading condition, students repeatedly read the materials in four consecutive reading periods. In a retrieval practice condition, students read the texts in one reading period and then practiced retrieving their knowledge by freely recalling the material. Students in this condition wrote down as much of the material as they could recall in three consecutive retrieval practice periods. In this particular experiment, the students did not reread the texts or receive feedback, and the two conditions were matched on the total time they spent studying the material. Figure 70.4A shows the effects of these two study conditions on a final free recall test. When the final test occurred a few minutes after the initial learning period, the repeated reading group performed slightly better than the retrieval practice group. However, when the final test instead occurred 1 week after the learning phase, the retrieval practice group far outperformed the repeated reading group. Repeated retrieval produced a 50% improvement on the delayed retention test.

A variety of theoretical ideas have been proposed to explain why retrieval practice enhances learning. One theory is that the process of retrieval affords semantic elaboration. This theory proposes that during retrieval, people mentally generate knowledge that is related to the content they are retrieving, and this elaboration adds details that help make target knowledge more distinctive and retrievable in the future (Carpenter, 2009, 2011). Other theorists have proposed that the process of retrieval directly promotes organizational or relational processing (Congleton & Rajaram, 2012; Zaromb & Roediger, 2010). A third theory is that practicing retrieval directly enhances the diagnosticity of potential retrieval cues, essentially by improving memory search processes (Karpicke et al., 2014; Lehman et al., 2014; Whiffen & Karpicke, 2017). Ultimately, in one way or another, all theories of retrieval practice come back to the

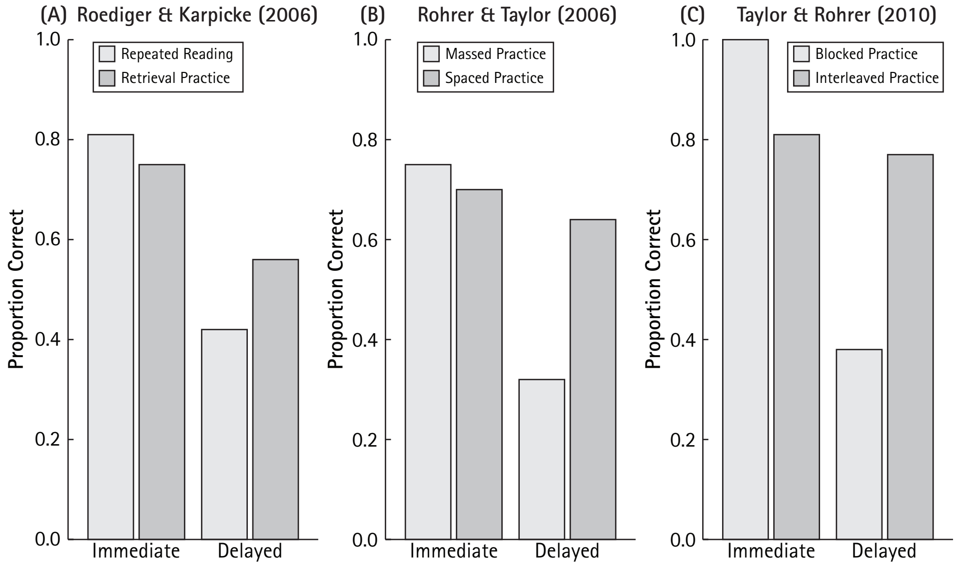

Figure 70.4 Three effective strategies that enhance long-term learning but may not be evident in the short term, reflecting a "desirable difficulty" pattern of results. (A) Results from Roediger and Karpicke (2006), showing that practicing retrieval enhanced retention on a delayed final test but not on a test shortly after learning. (B) Results from Rohrer and Taylor (2006), showing that spaced practice enhanced retention on a delayed test but not on an immediate test. (C) Results from Taylor and Rohrer (2010), showing large benefits of interleaved practice relative to blocked practice on a delayed test but not on an immediate test.

elements of effective learning. Retrieval practice enhances learning because it makes knowledge easier to reconstruct and use in future retrieval contexts.

There are many ways that instructors might use retrieval practice activities in their classrooms (Agarwal & Bain, 2019). Classroom quizzing or questioning activities, regardless of whether they are low- or high-tech (e.g., using personal response systems or "clickers"), are effective ways to implement retrieval practice in school settings (Roediger et al., 2011). Likewise, there are many ways students can practice retrieval when they study on their own. The most readily apparent way for Kim to incorporate retrieval practice into her studying would be to answer practice questions or practice recalling key terms, ideas, and concepts from her history class. But there are a few key things Kim could do to maximize the benefit of retrieval practice. One is to make sure she answers practice questions by retrieving knowledge, rather than simply reading and then looking up the answers to questions (as in an open-book quiz), which obviates retrieval. A second is to answer questions repeatedly, not just once. Many students will "self-test" when they study as a means of checking their knowledge, and they often stop after recalling an item or answering a question one time (Ariel & Karpicke, 2018; Karpicke, 2009). Even when a student can successfully retrieve knowledge once, repeatedly retrieving two or three additional times confers large benefits for learning (Grimaldi & Karpicke, 2014; Karpicke & Roediger, 2008). Finally, a third way Kim can ensure that she gets the most benefit from retrieval practice is to space her retrieval practice over time. That leads to the effective learning strategy discussed next.

Spaced Practice

The next two effective strategies, spaced practice and interleaved practice, complement each other. *Spacing* refers to distributing events over time rather than having the events occur close together in time, which is referred to as massed practice. Spaced practice can involve distributing topics at different points in time throughout a course, distributing study sessions over time, or distributing certain items such as questions or practice problems *within* a study session. In each scenario, people learn more by studying in a spaced fashion (studying a bit each day, or answering a question and then coming back to it later in a study session) than they do by studying in a massed fashion (cramming into a single large study session, or answering a question and then immediately repeating the answer over and over in one's head).

Spaced practice works in a wide range of settings and content areas, and an excellent example of the benefits of spaced practice in math education comes from research by Rohrer and Taylor (2006). In their experiments, college students learned math problems about permutations by solving a total of 10 practice problems. In a massed practice condition, students solved all 10 practice problems in a single session. In a spaced practice condition, students solved 5 practice problems in two separate sessions spaced 1 week apart. The two groups had exactly the same amount of practice; the only difference was whether the practice was crammed into one session or distributed across two sessions. The students then took a final test in which they solved a set of new permutation questions. Figure 70.4B shows that when the final test occurred 1 week after the last study session, there was a negligible difference between the two conditions, slightly favoring the massed practice condition. However, when retention was assessed 4 weeks after the last study session, there was a very large benefit of spaced practice, doubling performance on the test relative to massed practice (64% vs. 32%).

Like retrieval practice, a variety of theories have been proposed to explain the benefits of spaced practice. One theory is that spacing sessions, problems, or questions over time produces a more varied experience relative to when items or events happen close together in a massed fashion (for a critical examination, see Toppino & Gerbier, 2014). This idea, known as encoding variability, is directly tied to the idea of elaboration. Variable encoding is assumed to produce elaboration, increasing the likelihood that a learner will encounter effective retrieval cues in the future. A second related theory has to do with the "completeness" of the mental operations people perform when repetitions are massed or spaced. When students solve similar problems that are spaced apart, they may have to engage in more complete mental processes the second time than they would when the problems are massed together. Thus, learners benefit because the second study event is more complete and more effective with spaced practice (for a review of research on spaced practice, see Carpenter, 2017). Theories of spaced practice fit within the framework of the elements of effective learning. Spaced practice enhances elaborative encoding, which in turn makes knowledge more retrievable in a variety of retrieval contexts.

Spaced practice is effective *across* study sessions, by distributing sessions over time, and *within* sessions, by distributing specific problems or questions throughout a session. Instructors can do a great service to their students by providing review opportunities, such as problem sets or worksheets that cover previous topics, at regular spaced intervals. In her U.S. history class, Kim would benefit from briefly reviewing the key concepts she needs to know at spaced intervals, rather than cramming material into a lengthy and fatiguing study session right before an exam. Likewise, Josephine could space her study efforts by reviewing key ideas in her math class, such as the formulas she needs to know to solve particular problems, during a brief once-a-week review session, thereby spacing periodic reviews throughout the semester. One way to accomplish spaced practice within a learning session is to cycle through and repeat items or questions during a session so that students come back to particular items after they have spent time with other items, thereby introducing spaced practice. And one specific way to accomplish spacing within a session is to intermix different types of problems or questions within a session—a technique known as interleaving, which is discussed next.

Interleaved Practice

Interleaved practice refers to mixing together different types of items, questions, or problems, rather than practicing several of the same type of item in a row, which is known as blocked practice. A good example of interleaved practice comes from a study in which baseball players practiced hitting different types of pitches—fastballs, curveballs, and changeups—in a blocked fashion, where each type of pitch was practiced several times before moving on to other types, or in an interleaved condition, where different types of pitches were randomly mixed together. Players performed better in later games when they trained with interleaved practice rather than blocked practice (Hall et al., 1994). Interleaved practice enhances learning of a range of educational materials, but unfortunately, many education activities afford blocked practice—practicing the same skill consecutively or solving the same type of problem several times in a row. Interleaved practice represents a simple change to existing instructional practices that can create dramatic benefits for students.

Research by Taylor and Rohrer (2010) provides a demonstration of the power of interleaved practice. They had fourth-grade students complete math problems that involved identifying the total number of faces, corners, edges, and angles of prisms. Students in a blocked condition completed all of the practice problems about one feature of prisms (faces) before moving on to problems about the next feature (corners), whereas students in an interleaved condition completed the set of practice problems intermixed together. The results, shown in Figure 70.4C**,** depict a perfect example of a "desirable difficulty" pattern of results. At the end of the initial practice phase, students performed better in the blocked condition than they did in the interleaved condition. If learning were assessed only as performance during this initial practice phase, one would conclude that blocked practice was the more effective instructional method. However, the students were tested 1 day later on a new set of prism problems and the results, even after only a 1-day delay, were dramatic. Interleaved practice produced substantially better performance relative to blocked practice. Students in the interleaved group maintained their level of performance almost perfectly over the 1-day delay, whereas students in the blocked group showed dramatic forgetting, dropping from 100% to 38% after 1 day.

A prevailing theory is that interleaved practice works by promoting "discriminative contrast" (Birnbaum et al., 2013; Carvalho & Goldstone, 2019; Foster et al., 2019). When different types of items, problems, or questions are mixed together, the differences and contrasts among the problems are emphasized. In this way, interleaving promotes elaboration by highlighting the distinctiveness of different items, which in turn promotes successful retrieval in the future. Interleaving different types of items or problems is also a way to implement spaced practice because problems of the same type would be distributed throughout a study session or across a worksheet, rather than occurring all together in a massed fashion. Thus, interleaved practice introduces variability and promotes more completeness of the processing that occurs during each item or practice problem. Finally, interleaving practice problems requires learners to practice identifying the type of problem they are trying to solve, a skill that is obviated when several of the same type of problems occur together in a blocked fashion. Interleaving thus enhances the effectiveness of potential retrieval cues by improving the ability to interpret those cues and know what relevant knowledge to bring to mind in a given retrieval context.

In her math class, Josephine would learn far more by solving different types of problems interleaved together rather than practicing solving the same type of problem over and over. For example, an interleaved worksheet could include a set of problems where each one required calculating a different attribute (e.g., surface area or volume) of a variety of different objects (circles, rectangles, pyramids, cylinders, and so on). As noted earlier, unfortunately, blocked practice remains the norm in many educational settings, with students solving several of the same type of problem on a worksheet or homework activity (Rohrer et al., 2020). Although Josephine might spend time solving practice problems and diligently completing her homework, her efforts would be far more productive if the same problems were intermixed rather than blocked, providing Josephine the opportunity to practice interpreting and identifying different problem types and promoting distinctive encoding.

Elaborative Interrogation

The last strategy reviewed here has an ornate name—*elaborative interrogation*—but represents a simple activity: asking "why" questions that prompt learners to construct explanatory answers. Like retrieval practice, spacing, and interleaving, elaborative interrogation is a strategy that may represent a simple modification to existing practices. Many questions could be turned into or followed up with "why" questions to prompt elaborative explanations. It is important to note, as emphasized by Dunlosky et al. (2013), that the research base on elaborative interrogation has limitations. Most of the research on elaborative interrogation has used simple laboratory materials, and more work remains to be done in authentic classroom settings. Nevertheless, the strategy is included here because it is simple to implement and there is promising evidence of its effectiveness.

In a study by Wood et al. (1990), elementary and middle school students (aged 8–14 years) read nine stories that each described the life and habits of one animal (e.g., the Western spotted skunk, the blue whale, the emperor penguin, and others). In a control condition, students simply read each story, which described six facts about the animal. In an elaborative interrogation condition, students were asked to provide an explanation for each of the six facts contained in the story. For example, one fact was that the Western spotted skunk lives in a hole in the ground. Students in the elaborative interrogation condition had to generate an explanation about *why* this animal would live in a hole in the ground (e.g., it does so in order to protect itself and its family). On a final test, the students were given the animal facts they had learned and were asked to recall the names of the animals that each fact corresponded to. Students who had engaged in elaborative interrogation while studying retained more knowledge about the animals than did the students who had only read the materials.

Research on the theoretical underpinnings of elaborative interrogation is somewhat limited, but the technique bears some of the hallmarks of other elaborative strategies. "Why" questions would prompt learners to retrieve prior knowledge and generate explanations. Answering "why" questions might promote organizational processing, leading learners to notice the similarities among multiple terms or concepts, and would very likely promote distinctive encoding by requiring learners to come up with unique explanatory answers that embellish the specific items. Either way, "why" questions provide learners with the opportunity to add detail and enrich the encoding of the material they are learning.

"Why" questions would be simple to include in a variety of settings. Given the effectiveness of elaborative interrogation with younger learners, the strategy would work well in Mason's fourth-grade unit on weather phenomena. His teacher could begin a classroom discussion by asking the class to brainstorm and come up with reasons why different types of clouds would form in certain conditions. "Why" questions could be asked as follow-ups to factual questions, thus turning factual questions into opportunities for elaboration. For example, if Mason and his classmates were asked, "What types of dark clouds are likely to be in the air during a thunderstorm?" (cumulonimbus clouds), a follow-up question would be, "*Why* are cumulonimbus clouds related to thunderstorms?" (e.g., because they contain a lot of moisture, or because they are formed when moist air rises and condenses). Likewise, existing homework worksheets in Mason's class could be enhanced by including a few "why" questions to prompt elaboration. Elaborative interrogation is a promising learning strategy because of its simplicity: "Why" questions can be incorporated into a wide range of instructional situations.

Conclusion

This chapter has proposed that there are indeed principles of learning, universal truths at work in all learning and memory scenarios. Effective learning happens when people have retrieval cues available that permit them to express their knowledge, when those retrieval cues are diagnostic of target knowledge, and when elaborative study methods have prepared learners to use potential retrieval cues by promoting organization and distinctiveness during initial study. Principles of learning are at work in laboratory experiments with simple materials, in studies that involve more complex problem-solving, and in actual classroom settings. The elements of effective learning represent a framework that outlines what learners must do to coordinate their actions in a dynamic environment—when they are asked to recall facts, answer questions, make inferences, solve problems, and apply their knowledge in new situations.

As a final note, this chapter described effective learning strategies individually, but the strategies are not mutually exclusive. Students and teachers could combine strategies that involve explaining, retrieving, spacing, and interleaving in a variety of ways. However, very little research has examined different possible combinations of strategies in order to identify the most effective ways to combine multiple strategies (O'Day & Karpicke, 2021). This represents an exciting area awaiting more research and exploration that could lead to the creation of new pedagogical techniques grounded on the elements of effective learning.

Acknowledgments

The writing of this chapter was supported in part by grant 220020483 from the James S. McDonnell Foundation. The opinions expressed are those of the authors and do not represent the views of the foundation. We thank Michelle Coverdale, Stephanie Karpicke, and Amedee Martella for useful discussion and feedback.

References

Agarwal, P. K., & Bain, P. M. (2019). *Powerful teaching: Unleash the science of learning*. Jossey-Bass.

Ariel, R., & Karpicke, J. D. (2018). Improving self-regulated learning with a retrieval practice intervention. *Journal of Experimental Psychology: Applied*, *24*(1), 43–56. doi:10.1037/xap0000133

Birnbaum, M. S., Kornell, N., Bjork, E. L., & Bjork, R. A. (2013). Why interleaving enhances inductive learning: The roles of discrimination and retrieval. *Memory & Cognition*, *41*(3), 392–402. doi:10.3758/s13421-012-0272-7

Bjork, R. A. (1999). Assessing our own competence: Heuristics and illusions. In D. Gopher & A. Koriat (Eds.), *Attention and performance* (pp. 435–459). MIT Press.

Blasiman, R. N., Dunlosky, J., & Rawson, K. A. (2017). The what, how much, and when of study strategies: Comparing intended versus actual study behaviour. *Memory*, *25*(6), 784–792. doi:10.1080/09658211.2016.1221974

Callender, A. A., & McDaniel, M. A. (2009). The limited benefits of rereading educational texts. *Contemporary Educational Psychology*, *34*(1), 30–41. doi:10.1016/j.cedpsych.2008.07.001

Carpenter, S. K. (2009). Cue strength as a moderator of the testing effect: The benefits of elaborative retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*(6), 1563–1569. doi:10.1037/a0017021

Carpenter, S. K. (2011). Semantic information activated during retrieval contributes to later retention: Support for the mediator effectiveness hypothesis of the testing effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*(6), 1547–1552. doi:10.1037/a0024140

Carpenter, S. K. (2017). Spacing effects on learning and memory. In J. T. Wixted (Ed.), *Learning and memory: A comprehensive reference: Vol. 2. Cognitive psychology of memory* (pp. 465–485). Academic Press.

Carvalho, P. F., & Goldstone, R. L. (2019). When does interleaving practice improve learning? In J. Dunlosky & K. A. Rawson (Eds.), *The Cambridge handbook of cognition and education* (pp. 411–436). Cambridge University Press.

Congleton, A., & Rajaram, S. (2012). The origin of the interaction between learning method and delay in the testing effect: The roles of processing and conceptual retrieval organization. *Memory & Cognition*, *40*(4), 528–539. doi:10.3758/s13421-011-0168-y

Craik, F. I. M. (2002). Levels of processing: Past, present . . . and future? *Memory*, *10*(5–6), 305–318. doi:10.1080/09658210244000135

Craik, F. I. M., & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning & Verbal Behavior*, *11*(6), 671–684. doi:10.1016/s0022-5371(72)80001-x

Craik, F. I. M., & Tulving, E. (1975). Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology: General*, *104*(3), 268–294. doi:10.1037/0096-3445.104.3.268

Craik, F. I. M., & Watkins, M. J. (1973). The role of rehearsal in short-term memory. *Journal of Verbal Learning & Verbal Behavior, 12*(6), 599–607. doi:10.1016/s0022-5371(73)80039-8

Duncker, K. (1945). *On problem-solving* (Psychological Monographs No. 270). American Psychological Association.

Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J., & Willingham, D. T. (2013). Improving students' learning with effective learning techniques: Promising directions from cognitive and educational psychology. *Psychological Science in the Public Interest*, *14*(1), 4–58. doi:10.1177/1529100612453266

Fiorella, L., & Mayer, R. E. (2014). Role of expectations and explanations in learning by teaching. *Contemporary Educational Psychology*, *39*(2), 75–85. doi:10.1016/j.cedpsych.2014.01.001

Fiorella, L., & Mayer, R. E. (2015). *Learning as a generative activity: Eight learning strategies that promote understanding*. Cambridge University Press.

Foster, N. L., Mueller, M. L., Was, C., Rawson, K. A., & Dunlosky, J. (2019). Why does interleaving improve math learning? The contributions of discriminative contrast and distributed practice. *Memory & Cognition*, *47*(6), 1088–1101. doi:10.3758/s13421-019-00918-4

Gentner, D., & Maravilla, F. (2018). Analogical reasoning. In L. J. Ball & V. A. Thompson (Eds.), *The Routledge international handbook series* (pp. 186–203). Routledge.

Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, *12*(3), 306–355. doi:10.1016/0010-0285(80)90013-4

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, *15*(1), 1–38. doi:10.1016/0010-0285(83)90002-6

Grimaldi, P. J., & Karpicke, J. D. (2014). Guided retrieval practice of educational materials using automated scoring. *Journal of Educational Psychology*, *106*(1), 58–68. doi:10.1037/a0033208

Hall, K. G., Domingues, D. A., & Cavazos, R. (1994). Contextual interference effects with skilled baseball players. *Perceptual and Motor Skills*, *78*(3, Pt. 1), 835–841. doi:10.2466/pms.1994.78.3.835

Hirst, W., Phelps, E. A., Meksin, R., Vaidya, C. J., Johnson, M. K., Mitchell, K. J., Buckner, R. L., Budson, A. E., Gabrieli, J. D., Lustig, C., Mather, M., Ochsner, K. N., Schacter, D., Simons, J. S., Lyle, K. B., Cuc, A. F., & Olsson, A. (2015). A ten-year follow-up of a study of memory for the attack of September 11, 2001: Flashbulb memories and memories for flashbulb events. *Journal of Experimental Psychology: General*, *144*(3), 604–623. doi:10.1037/xge0000055

Hunt, R. R., & McDaniel, M. A. (1993). The enigma of organization and distinctiveness. *Journal of Memory and Language*, *32*(4), 421–445. doi:10.1006/jmla.1993.1023

Hunt, R. R., & Smith, R. E. (1996). Accessing the particular from the general: The power of distinctiveness in the context of organization. *Memory & Cognition*, *24*(2), 217–225. doi:10.3758/bf03200882

Kahana, M. J. (2017). Memory search. In J. T. Wixted (Ed.), *Learning and memory: A comprehensive reference: Vol. 2. Cognitive psychology of memory* (pp. 181–200). Academic Press.

Karpicke, J. D. (2009). Metacognitive control and strategy selection: Deciding to practice retrieval during learning. *Journal of Experimental Psychology: General*, *138*(4), 469–486. doi:10.1037/a0017341

Karpicke, J. D. (2017). Retrieval-based learning: A decade of progress. In J. T. Wixted (Ed.), *Learning and memory: A comprehensive reference: Vol. 2. Cognitive psychology of memory* (pp. 487–514). Academic Press.

Karpicke, J. D., Butler, A. C., & Roediger, H. L. (2009). Metacognitive strategies in student learning: Do students practice retrieval when they study on their own? *Memory*, *17*(4), 471–479. doi:10.1080/09658210802647009

Karpicke, J. D., Lehman, M., & Aue, W. R. (2014). Retrieval-based learning: An episodic context account. Psychology of Learning and Motivation, *61*, 237–284.

Karpicke, J. D., & Roediger, H. L. (2008). The critical importance of retrieval for learning. *Science*, *319*(5865), 966–968. doi:10.1126/science.1152408

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. University of Chicago Press.

Lehman, M., Smith, M. A., & Karpicke, J. D. (2014). Toward an episodic context account of retrieval-based learning: Dissociating retrieval practice and elaboration. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*(6), 1787–1794. doi:10.1037/xlm0000012

Mäntylä, T. (1986). Optimizing cue effectiveness: Recall of 500 and 600 incidentally learned words. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*(1), 66–71. doi:10.1037/0278-7393.12.1.66

McDaniel, M. A., & Einstein, G. O. (1989). Material-appropriate processing: A contextualist approach to reading and studying strategies. *Educational Psychology Review*, *1*(2), 113–145. doi:10.1007/BF01326639

McDaniel, M. A., & Einstein, G. O. (2005). Material appropriate difficulty: A framework for determining when difficulty is desirable for improving learning. In A. F. Healy (Ed.), *Experimental cognitive psychology and its applications* (pp. 73–85). American Psychological Association.

Nairne, J. S. (2002). The myth of the encoding-retrieval match. *Memory*, *10*(5–6), 389–395. doi:10.1080/09658210244000216

Nairne, J. S. (2006). Modeling distinctiveness: Implications for general memory theory. In R. R. Hunt & J. B. Worthen (Eds.), *Distinctiveness and memory* (pp. 27–46). Oxford University Press.

Nunes, L. D., & Karpicke, J. D. (2015). Retrieval-based learning: Research at the interface between cognitive science and education. In R. A. Scott & S. M. Kosslyn (Eds.), *Emerging trends in the social and behavioral sciences* (pp. 1–16). Wiley.

O'Day, G. M., & Karpicke, J. D. (2021). Comparing and combining retrieval practice and concept mapping. *Journal of Educational Psychology*, *113*(5), 986–997. doi:10.1037/edu0000486

Pashler, H., Bain, P. M., Bottge, B. A., Graesser, A., Koedinger, K., McDaniel, M., & Metcalfe, J. (2007). *Organizing instruction and study to improve student learning: A practice guide* (NCER 2007-2004). Institute of Education Sciences.

Rawson, K. A., & Kintsch, W. (2005). Rereading effects depend on time of test. *Journal of Educational Psychology*, *97*(1), 70–80. doi:10.1037/0022-0663.97.1.70

Roediger, H. L. (1980). Memory metaphors in cognitive psychology. *Memory & Cognition*, *8*(3), 231–246. doi:10.3758/bf03197611

Roediger, H. L., III. (2008). Relativity of remembering: Why the laws of memory vanished. *Annual Review of Psychology*, *59*, 225–254. doi:10.1146/annurev.psych.57.102904.190139

Roediger, H. L., III, Agarwal, P. K., McDaniel, M. A., & McDermott, K. B. (2011). Test-enhanced learning in the classroom: Long-term improvements from quizzing. *Journal of Experimental Psychology: Applied*, *17*(4), 382–395. doi:10.1037/a0026252

Roediger, H. L., & Karpicke, J. D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science*, *17*(3), 249–255. doi:10.1111/j.1467-9280.2006.01693.x

Roediger, H. L., III, Tekin, E., & Uner, O. (2017). Encoding-retrieval interactions. In J. T. Wixted (Ed.), *Learning and memory: A comprehensive reference: Vol. 2. Cognitive psychology of memory* (pp. 5–26). Academic Press.

Rohrer, D., Dedrick, R. F., Hartwig, M. K., & Cheung, C.-N. (2020). A randomized controlled trial of interleaved mathematics practice. *Journal of Educational Psychology*, *112*, 40–52. doi:10.1037/edu0000367

Rohrer, D., & Taylor, K. (2006). The effects of overlearning and distributed practise on the retention of mathematics knowledge. *Applied Cognitive Psychology*, 20(9), 1209–1224. doi:10.1002/acp.1266

Schwartz, D. L., Chase, C. C., Oppezzo, M. A., & Chin, D. B. (2011). Practicing versus inventing with contrasting cases: The effects of telling first on learning and transfer. *Journal of Educational Psychology*, *103*(4), 759–775. doi:10.1037/a0025140

Schwartz, D. L., & Goldstone, R. (2016). Learning as coordination: Cognitive psychology and education. In L. Corno & E. M. Anderman (Eds.), *Handbook of educational psychology* (3rd ed., pp. 61–75). Routledge.

Surprenant, A. M., & Neath, I. (2009). *Principles of memory*. Psychology Press.

Taylor, K., & Rohrer, D. (2010). The effects of interleaved practice. *Applied Cognitive Psychology*, *24*(6), 837–848. doi:10.1002/acp.1598

Thomson, D. M., & Tulving, E. (1970). Associative encoding and retrieval: Weak and strong cues. *Journal of Experimental Psychology*, *86*(2), 255–262. doi:10.1037/h0029997

Toppino, T. C., & Gerbier, E. (2014). About practice: Repetition, spacing, and abstraction. Psychology of Learning and Motivation, *60*, 113–189.

Tulving, E. (1983). *Elements of episodic memory*. Oxford University Press.

Tulving, E., & Madigan, S. A. (1970). Memory and verbal learning. *Annual Review of Psychology*, *21*, 437–484. doi:10.1146/annurev.ps.21.020170.002253

Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning & Verbal Behavior*, 5(4), 381–391. doi:10.1016/s0022-5371(66)80048-8

Tulving, E., & Thomson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review*, *80*(5), 352–373. doi:10.1037/h0020071

Watkins, O. C., & Watkins, M. J. (1975). Buildup of proactive inhibition as a cue-overload effect. *Journal of Experimental Psychology: Human Learning and Memory*, *1*(4), 442–452. doi:10.1037/0278-7393.1.4.442

Wechsler, D. (1963). Engrams, memory storage, and mnemonic coding. *American Psychologist*, *18*(3), 149–153. doi:10.1037/h0046456

Whiffen, J. W., & Karpicke, J. D. (2017). The role of episodic context in retrieval practice effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *43*(7), 1036–1046. doi:10.1037/xlm0000379

Wood, E., Pressley, M., & Winne, P. H. (1990). Elaborative interrogation effects on children's learning of factual content. *Journal of Educational Psychology*, *82*(4), 741–748. doi:10.1037/0022-0663.82.4.741

Zaromb, F. M., & Roediger, H. L., III. (2010). The testing effect in free recall is associated with enhanced organizational processes. *Memory & Cognition*, *38*(8), 995–1008. doi:10.3758/MC.38.8.995