Integrating COGNITIVE SCIENCE with INNOVATIVE TEACHING in STEM DISCIPLINES

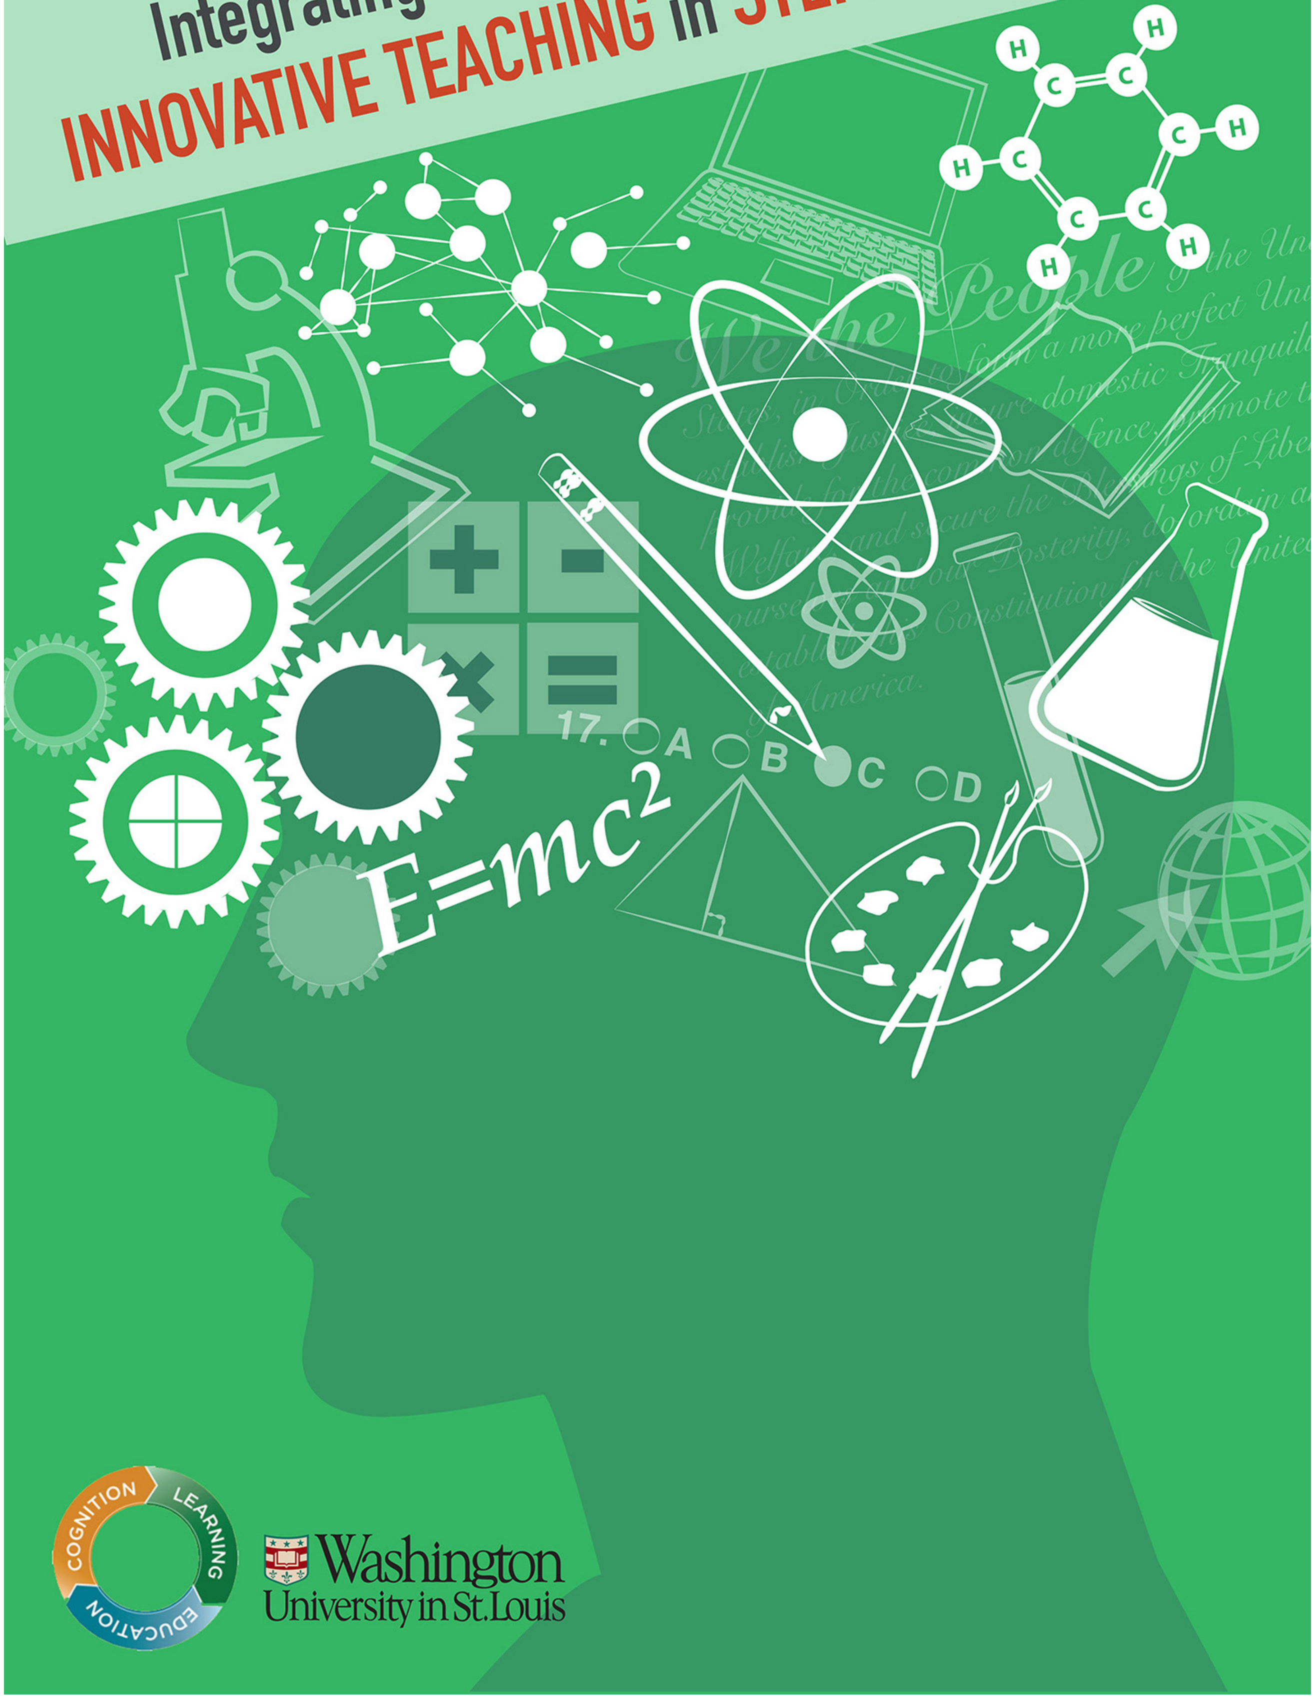

Integrating Cognitive Science with Innovative Teaching in STEM Disciplines

Edited by

Mark A. McDaniel, Regina F. Frey, Susan M. Fitzpatrick, & Henry L. Roediger, III

Washington University Libraries

St. Louis, Missouri

2014

1. The Increasing Importance of Learning How to Learn

doi:10.7936/K7QN64NR
Robert A. Bjork and Veronica X. Yan

University of California, Los Angeles

Correspondence:
Robert A. Bjork
Department of Psychology
University of California
Los Angeles CA 90095-1563
http://bjorklab.psych.ucla.edu/

Abstract

Increasingly, learning is happening outside of formal classroom instruction. As a consequence, learners need to make multiple decisions, such as what to study, when to study, and how to study, and computer-based technologies offer multiple options and opportunities for how to manage one's own learning. Knowing how to learn effectively has never been more important, not only during the years of schooling, but across one's lifetime—as careers change, new job skills are required, and hobbies and interests develop and change. Recent research suggests, however, that we are often prone to both mis-assessing and mis-managing our own learning. In this chapter we summarize the evidence that intuitions and standard practices are often unreliable guides to optimizing one's learning and that there exists the potential for learners and instructors alike to make self-regulated and teacher-regulated learning more efficient and effective.

Introduction

For understandable reasons, instructors at all levels are interested in how they should teach their students, and considerable research attention has focused on how lectures and assignments can be structured to enhance students' learning and comprehension. A topic that is arguably even more important, however, namely, what learners can do to enhance their own learning has received much less attention. Knowing how to manage and assess one's own learning has always been important, but learning how to learn has become increasingly crucial. In a world that is not only complex and rapidly changing, but also characterized by technologies, such as online courses and classrooms, podcasts, and the myriad learning opportunities afforded by the Internet, classroom-type learning is being pushed more and more outside of the classroom.

Knowing how to learn effectively outside the classroom becomes especially important during and after college. During the K-12 years, most of us sat, or sit, in class listening to teachers or engaging in class activities for up to seven hours a day and were, or are, assigned regular homework. In college, however, learning takes place largely outside of the classroom and is mostly in our own hands as learners. Furthermore, beyond the years of formal education, and increasingly across the lifespan, learning is almost exclusively the responsibility of the individual learner. Thus, not only is it important to be concerned with how instructors should teach, but it is also critical that we teach our students what Bjork, Dunlosky, and Kornell (2013) have recently labeled the "ultimate survival tool," namely, how to learn.

Components of Becoming Sophisticated as a Learner

Becoming a metacognitively sophisticated learner is not a trivial matter, because an efficient and effective learner has to be able to both monitor and control his or her own learning effectively, which can require overcoming certain intuitions and impressions. Monitoring one's own memory accurately involves knowing whether information or procedures have been learned to a degree that will support their later recall and transfer when needed, which could be in the context of an examination, or in the context of a job. It is important not only that additional study time be allocated when such learning goals have not yet been achieved, but also that study time not be wasted when those goals have already been achieved. While this sounds simple enough, recent research has shown that learners often are not accurate in monitoring their own learning.

That such monitoring can be faulty is illustrated by real-world experiences we have all had. We can, for example, imagine times where we have gone into an exam lacking confidence but then scored well, or, conversely, gone into an exam full of confidence and then scored poorly. Similarly, based on watching somebody else execute some to-be-learned procedure, we might experience a sense of complete understanding, but then, later, find out that we have no idea what to do next when it is required that we execute the procedure. All of these examples are instances of imperfect monitoring.

Achieving accurate monitoring is, however, only part of the battle. Even when we are able to monitor our own learning accurately, we must then understand how to control our learning activities effectively. If a learner identifies something as requiring more study, the next step is to know how to go about gaining that knowledge. Study strategies are not made equal, however, and knowing how to go about scheduling one's own learning, both effectively and efficiently, is critical, particularly when time is limited, as it often is in today's busy lives.

Why Aren't We Already Effective Learners?

Every one of us is a lifetime learner. We have been learning since birth and every day thereafter, through our schooling years and beyond. From that standpoint, one might expect that we would be educated by the "trials and errors of everyday living and learning" (Bjork, 2011) and become experts at managing the conditions of our own learning, but that appears, surprisingly, not to be the case. Instead, research findings have demonstrated that we are susceptible to both mis-assessing and mis-managing our own learning.

A major reason that we can be fooled as to whether we have learned and how we should learn is that one's current performance and the subjective ease of processing are often poor indicators of long-term retention and transfer. In fact, as we describe below, there are conditions of instruction that make performance improve rapidly and thus may make it appear and feel as though we are learning, but which do not support long-term learning. Conversely, there are many conditions of instruction that appear to create difficulties for the learner and slow the rate of apparent learning but enhance long-term learning. These latter conditions of instruction may be considered "desirable difficulties" (Bjork, 1994) in that they engage learners in deeper, more elaborate, and more effortful processing. The mismatch between current performance and long-term learning has huge implications for how we as learners, as instructors, and as parents assess and guide learning.

In this chapter we focus on only a few of the different conditions of instruction that can introduce desirable difficulties for learners, namely, distributing practice, increasing contextual interference, and engaging in test-induced retrieval practice. Each of these manipulations, in our view, has important implications for the learning of science.

What Constitute Effective Study Strategies and Do Learners Use them?

Distributing practice.

True long-term learning requires repeated studying of information. The way in which repeated study opportunities are distributed, however, makes a large difference in whether information will be retained. If you are going to read a passage twice, should you study it twice consecutively or read it once and then wait before reading it a second time? Many learners may feel tempted to restudy the passage immediately after the first study in an effort to gain clarity about things that were not clear during the first reading. An immediate restudy can also feel easier and can convey a sense of fluency (which, unfortunately, can be confused with comprehension and understanding), whereas, when a gap is introduced between study and re-study, learners can sense that they have forgotten information in the interval, making the restudy session feel less productive. Zechmeister and Shaughnessy (1980), for example, found that participants judged information to be less well learned after spaced repetitions than after massed repetitions. Figure 1 displays schematically an important pattern of results that has emerged from decades of research on the effects of spacing repeated study opportunities. As shown in the Figure, performance on tests administered after a very short delay often show a benefit of massed practice, as in the left side of Figure 1, whereas—and often contrary to learners' judgments—spaced practice enhances long-term learning, often substantially, as revealed by superior retention on delayed tests (e.g., Estes, 1955).

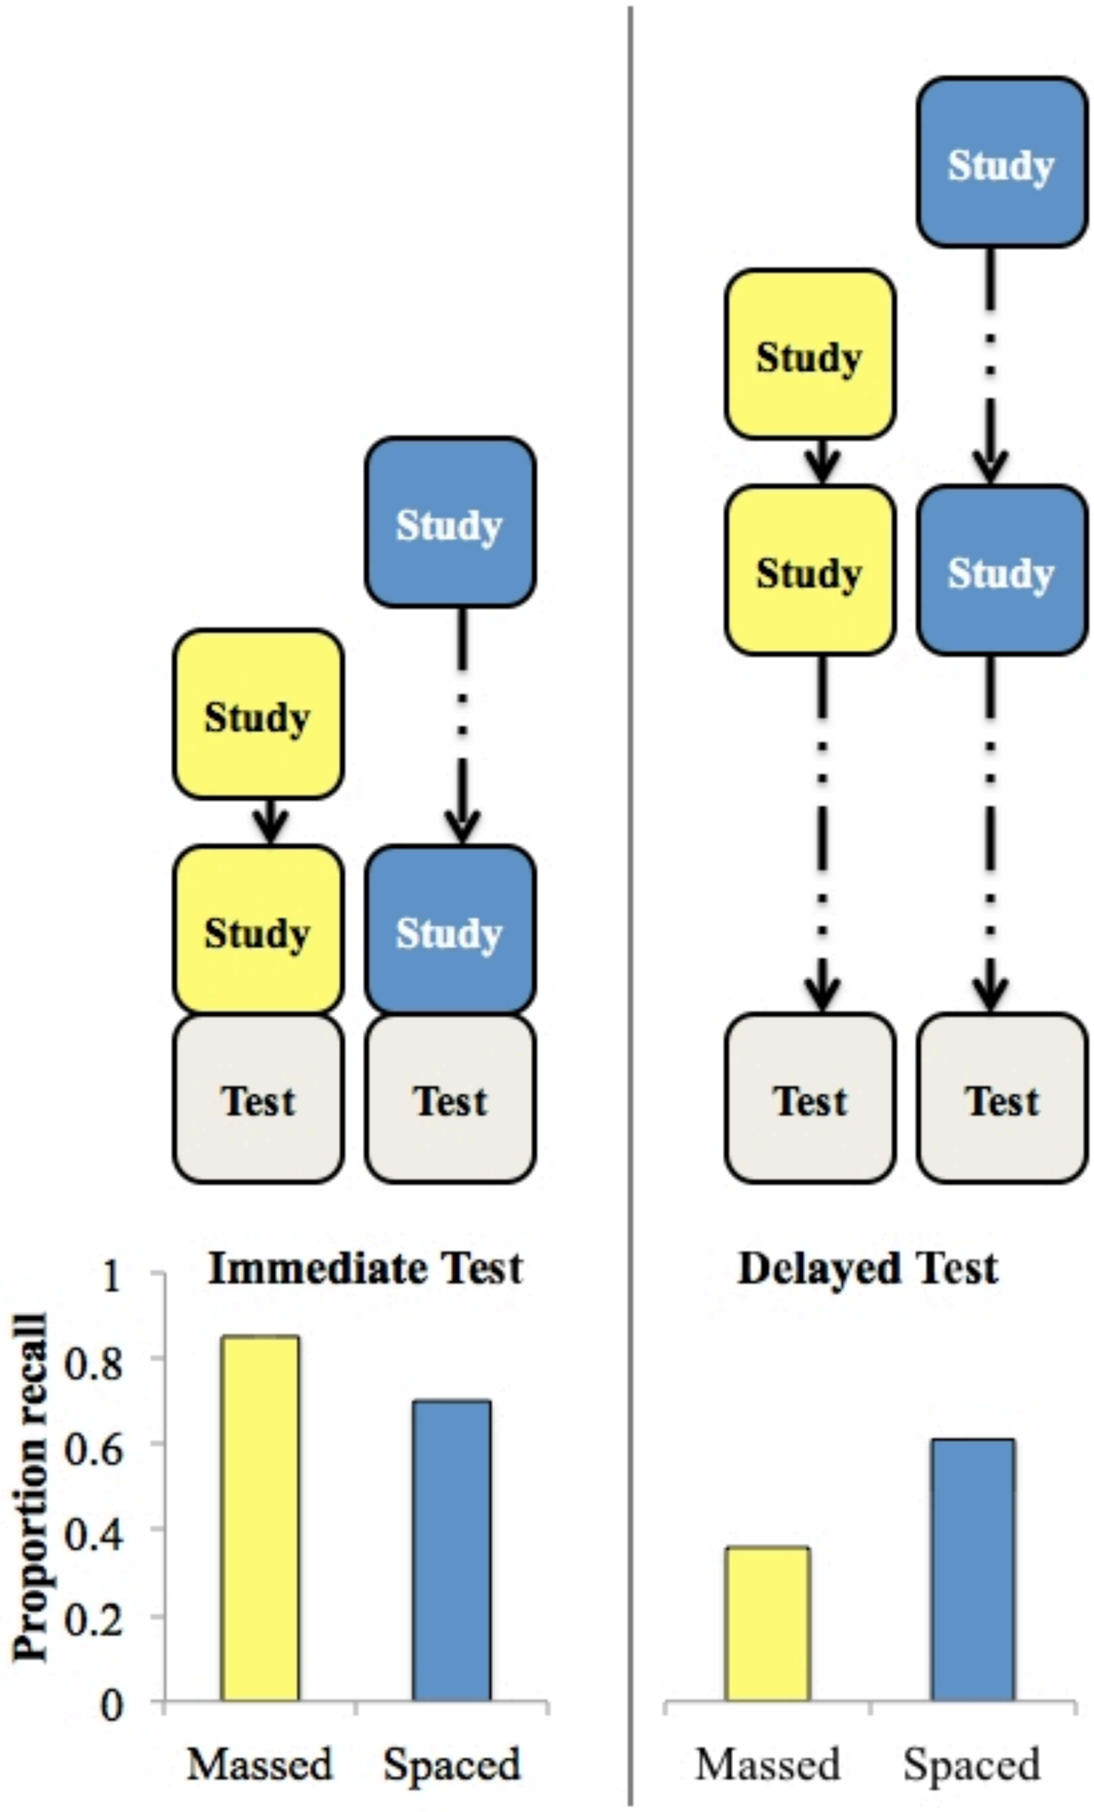

Figure 1. *The typical design of a spacing experiment and hypothetical results showing the typical effects of spaced and massed practice: Massing study may lead to better performance in the short-term, but spaced study yields better long-tern retention.*

The "spacing effect," which refers to long-term benefits of spacing, rather than massing, repeated study sessions, is one of the most robust findings from the entire history of experimental psychology (reported as early as Ebbinghaus, 1885/1964), and has been repeatedly demonstrated across a number of time scales, from seconds (e.g., Peterson & Peterson, 1959) to months (e.g., Bahrick, Bahrick, Bahrick & Bahrick, 1993), and across a variety of domains, from verbal learning (see Cepeda, Pashler, Vul, Wixted, & Rohrer, 2006, for a review) to advertisements (Appleton-Knapp, Bjork & Wickens, 2005) to motor skills learning (e.g., Shea, Lai, Black & Clark, 2000). It may feel more difficult to restudy information after an interval because access to information stored earlier can be lost across that interval (i.e., forgetting happens), but it is exactly this loss of access that can make spaced restudy more effective than immediate restudy. Immediately after an initial study, the to-be-learned information is highly accessible. Restudying the information at this time has a relatively small effect on overall learning, because such restudying does not engage any active retrieval processes. On the other hand, if learners are allowed to forget the information in between study and restudy, that restudy episode triggers retrieval of the first episode and thus constitutes a much more potent learning event (see the study-phase-retrieval interpretation of spacing effects; e.g., Appleton-Knapp, Bjork, & Wickens, 2005; Thios & D'Agostino, 1976).

More recently, this theory has been re-conceptualized in the framework of "reminding" (Benjamin & Ross, 2011): Restudying reminds learners of initial-study experiences, and the longer the lag, or the more dissimilar the restudy experience is from the initial study experience, the more potent the reminding as a learning event (with a caveat: If the restudy is delayed or dissimilar to the point that it does not trigger reminding, then learning does not profit). One might in fact state that forgetting leads to learning (given a restudy opportunity), rather than undoes learning. Said differently, conditions that lead to loss of access, such as increasing the length or difficulty of intervening activity (see, e.g., Bjork & Allen, 1970) or introducing a change of context (see, e.g., Smith, Glenberg, & Bjork, 1978), are also the conditions can lead to more effective relearning.

Despite the overwhelming evidence in favor of spacing, it appears that students do not appreciate the long-term benefits of spacing. One reason, perhaps, is that cramming (massed practice) right before an exam can produce good performance on the exam, as illustrated in the left side of Figure 1. The fact that such cramming will produce poor long-term retention is not, typically, something students become aware of, whereas they do experience and become aware that cramming can often, if not always, produce good exam scores. Hartwig and Dunlosky (2012), for example, found that 66 percent of surveyed students report cramming the night before exams, and Kornell and Bjork (2007) in a survey of 472 undergraduate students found that 64 percent of students surveyed said they did not restudy information if they feel like they had learned it.

Such survey findings provide some evidence that students fail to appreciate the benefits of spacing, rather than massing, repeated study sessions. Additional support comes from experiments in which learners are asked to make study decisions. In several experiments in which participants were asked whether they wanted to study individual vocabulary words "sooner," which meant during the first of two restudy periods, or "later," during a second restudy period, Cohen, Yan, Halamish, & Bjork (2013) found that participants wanted to study high-value or difficult items sooner rather than later. Importantly, the authors' experimental design made it clear to participants that the retention interval from either of the two restudy periods to the retention test on the items re-studied during that period would be the same. When the final retention interval is not controlled between the items that are studied "sooner" and "later", learners might choose to restudy valuable or more difficult items later, not because they understand the benefits of spacing, but, instead, because they want to place the study of those items closer to the final test. Cohen *et al*.'s findings suggests that participants do not appreciate that longer spacing is more beneficial than is shorter spacing for long-term retention.

Interleaving, rather than blocking, practice on different tasks.

When one really wants to learn complex concepts or categories, rather than just memorizing simple words or facts, one common piece of advice is to sit down and really focus on it. This is, for example, how textbooks are often organized: We immerse ourselves in studying and practicing one concept so that we can make sure we understand it fully before moving onto the next. Why do we not mix our study up, alternating back and forth between different concepts? The answer seems obvious: That would be confusing!

In the face of this very compelling intuition, however, research has demonstrated that providing "contextual interference" (Battig, 1972)—that is, arranging the instruction or practice of separate to-be-learned tasks or topics in a way the maximizes, rather than minimizes, the possible interference between those tasks—can enhance long-term retention and transfer. Initially, dating back to a classic paper by Shea and Morgan (1979), the benefits of interleaving were demonstrated in the learning of motor tasks, such as interleaving, rather than blocking, practice trials on the three different types of badminton serves (for a review, see Lee, 2012). More recently, however, research has demonstrated the benefits of interleaving extend to learning categories and concepts from examples (e.g., Kornell & Bjork, 2008) and to the learning of procedures, such calculating the volumes of different solids (e.g., Rohrer & Taylor, 2007).Kornell and Bjork (2008) investigated whether the learning of individual artists' styles was facilitated by presenting examples of a given artist's paintings in succession or presenting those paintings interleaved among examples of paintings by the other to-be-learned artists. More specifically, participants were shown six paintings by each of twelve artists. The paintings were shown one by one, with the paintings of a given artist shown either in immediate succession (the blocked condition) or intermixed among the paintings of other artists (the interleaved condition). On the final test, participants were shown new paintings by the studied artists and had to identify which artist painted each new painting.

Thus, participants could not simply rely upon memory for any individual painting, but had to have abstracted some idea of each artist's style.

Despite the intuitions of the researchers—and, as it turned out, the participants themselves—the interleaved condition, not the blocked condition, resulted in participants being better able to identify the artist responsible for a given new painting on the final test. Even though the participants were given feedback after each trial of the test, they overwhelmingly—when asked whether blocked or interleaved study helped them learn better—said that blocking was better. Kornell and Bjork speculated that whereas blocking exemplars of a given category might help learners notice commonalities within a category, interleaving juxtaposes exemplars of different categories, highlighting the differences between categories, which is apparently the more important consideration at the time of the final test. Subsequently, Kang and Pasher (2012) obtained evidence favoring the idea that interleaving enhances, in their term, "discriminative contrasts."Rohrer and Taylor (2007) have also demonstrated the benefit of interleaving over blocking with mathematics learning. College students learned to compute the volumes of different geometric solids in two practice sessions (spaced one week apart). In each of these sessions, students either practiced four problems of each of four different geometric solids in a blocked or interleaved manner. Throughout these two practice sessions, those in the blocked practice condition were consistently more accurate than those in the interleaved practice condition. As illustrated in Figure 2, however, this pattern was reversed on the criterion test one week later: Those who had practiced the problems interleaved had learned the formulas significantly better than those who practiced the problems blocked. Rohrer and Taylor argued that students who had undergone interleaved practice not only learned how to solve each kind of problem, but also learned to discriminate which formula was appropriate for each problem, a claim supported by the fact that the majority of errors made by the blockers took the form of applying the incorrect formula.

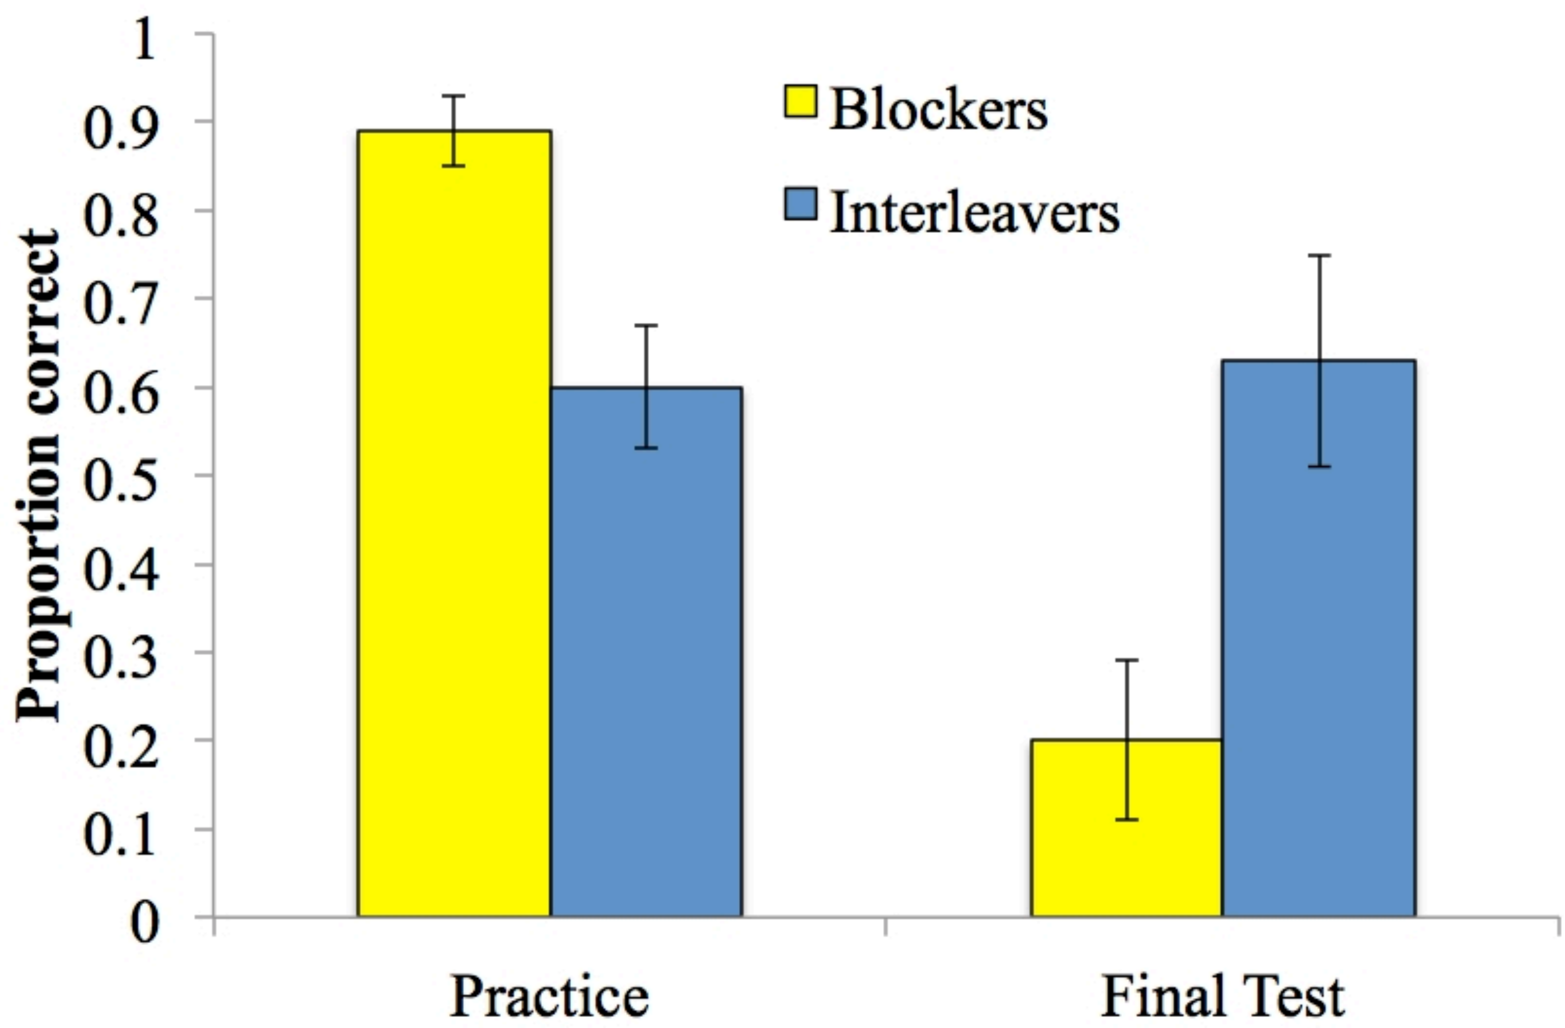

Figure 2. *The accuracy in solving math problems at the end of practice or on a delayed criterion test by participants who learned and practiced volume formulas blocked by type of solid or intermixed (from Rohrer and Taylor, 2007). Error bars represent one standard error of the mean.*

It is easy to understand how learners might underappreciate the benefits of interleaving study: Not only does blocked practice feel easier, it can also produce better performance during the learning process. In fact, it appears, based on recent research by Tauber, Dunlosky, Rawson, Wahlheim, and Jacoby (2012), that prior to any laboratory testing, learners hold an intuition that blocking practice is optimal. Tauber *et al*. gave participants the task of learning 12 bird families from examples of each family but let the participants make their own study decisions as to what example of what family they would like to see next. Overwhelmingly, the participants chose to block—that is, to see another example of a given family, rather than to see an example of a different family.

Questions still remain, however, as to how universally applicable interleaving benefits are: There are fewer studies on interleaving than there are on spacing, and existing literature has focused particularly on the interleaving of related concepts. There is now some evidence, using artificial rule-based categories, that the interleaving benefit is eliminated in cases where to-be-learned categories are relatively dissimilar and thus, highlighting commonalities within a category may be relatively more beneficial for learning (Carvalho & Goldstone, 2014). Also, it is less clear how the mechanisms of interleaving will work with unrelated materials. Should a student, for instance, interleave study of history, psychology, chemistry and literature? If the benefit of interleaving arises from enhanced discrimination between topics, it may not make sense to juxtapose completely unrelated topics, as there is unlikely anything in chemistry that will enhance the learning of history. On the other hand, given that interleaving introduces spacing—i.e., time elapses between when one first studies a topic and next returns to it—it may be that even in the absence of the useful contrasts between unrelated topics, interleaving the study of unrelated topics may still be advantageous. Recent findings reported by Birnbaum, Kornell, Bjork, and Bjork (2013) support the notion that when learning related topics, the opportunity to contrast these different topics and the spacing between examples within a given topic both contribute to the benefit of interleaving over blocking, so it may well be that interleaving unrelated as well as related topics is a good study strategy.

Finally, one practical benefit of interleaving relative to conventional spacing—given the reality of the time constraints in everyday life—is that implementing interleaving does not require a longer total period of instruction or time on task. That is, whereas spacing takes more total elapsed time than massing, interleaving takes the same amount of time as blocking. In both blocked and interleaved schedules, the same amount of information is studied across the same period of time; simply rearranging the order in which that information is studied can yield greater long-term learning.

Using tests as learning events.

By the end of their college years (and often to their chagrin), students have become experienced test-takers. From pop quizzes, to midterms, to final exams, tests in education are overwhelmingly used as high-stakes tests of assessment. The utility of tests, however, reaches far beyond simple assessment. Rather, tests are potent learning tools, and confer learning benefits that are often substantially greater than the benefits the result from re-reading. The results of an experiment by Roediger and Karpicke (2006) provide a good example of the point. Participants were asked to learn the content of a prose passage (either about "The Sun" or "Sea Otters"). Some participants studied this passage four times; others studied it three times and then took one free-recall test; and still others studied the passage only once and took three successive free-recall tests. Importantly, when participants took these free-recall tests, no feedback was provided. The results of the final criterion test, which was administered either five minutes or one week later, are displayed in Figure 3. On an immediate criterion test, those who studied more, recalled more. The pattern completely reversed, however, after one week: Participants who received more tests during the initial study session recalled more "idea units" from the passage.

Again, though, as in the case of interleaved versus blocked practice, the metacognitive judgments of the learners did not reflect the observed pattern. When asked at the end of the study phase how much they would remember in one week's time, those who had studied more gave higher ratings than did those who were tested more. Similarly, in Kornell and Bjork's (2007) survey of undergraduates' study habits, a majority 68 percent of the undergraduates reported using tests to assess their own learning, whereas only 18 percent indicated that they would learn more through tests than through rereading. Even more worryingly, nine percent of the respondents said they did not use testing for any reason.

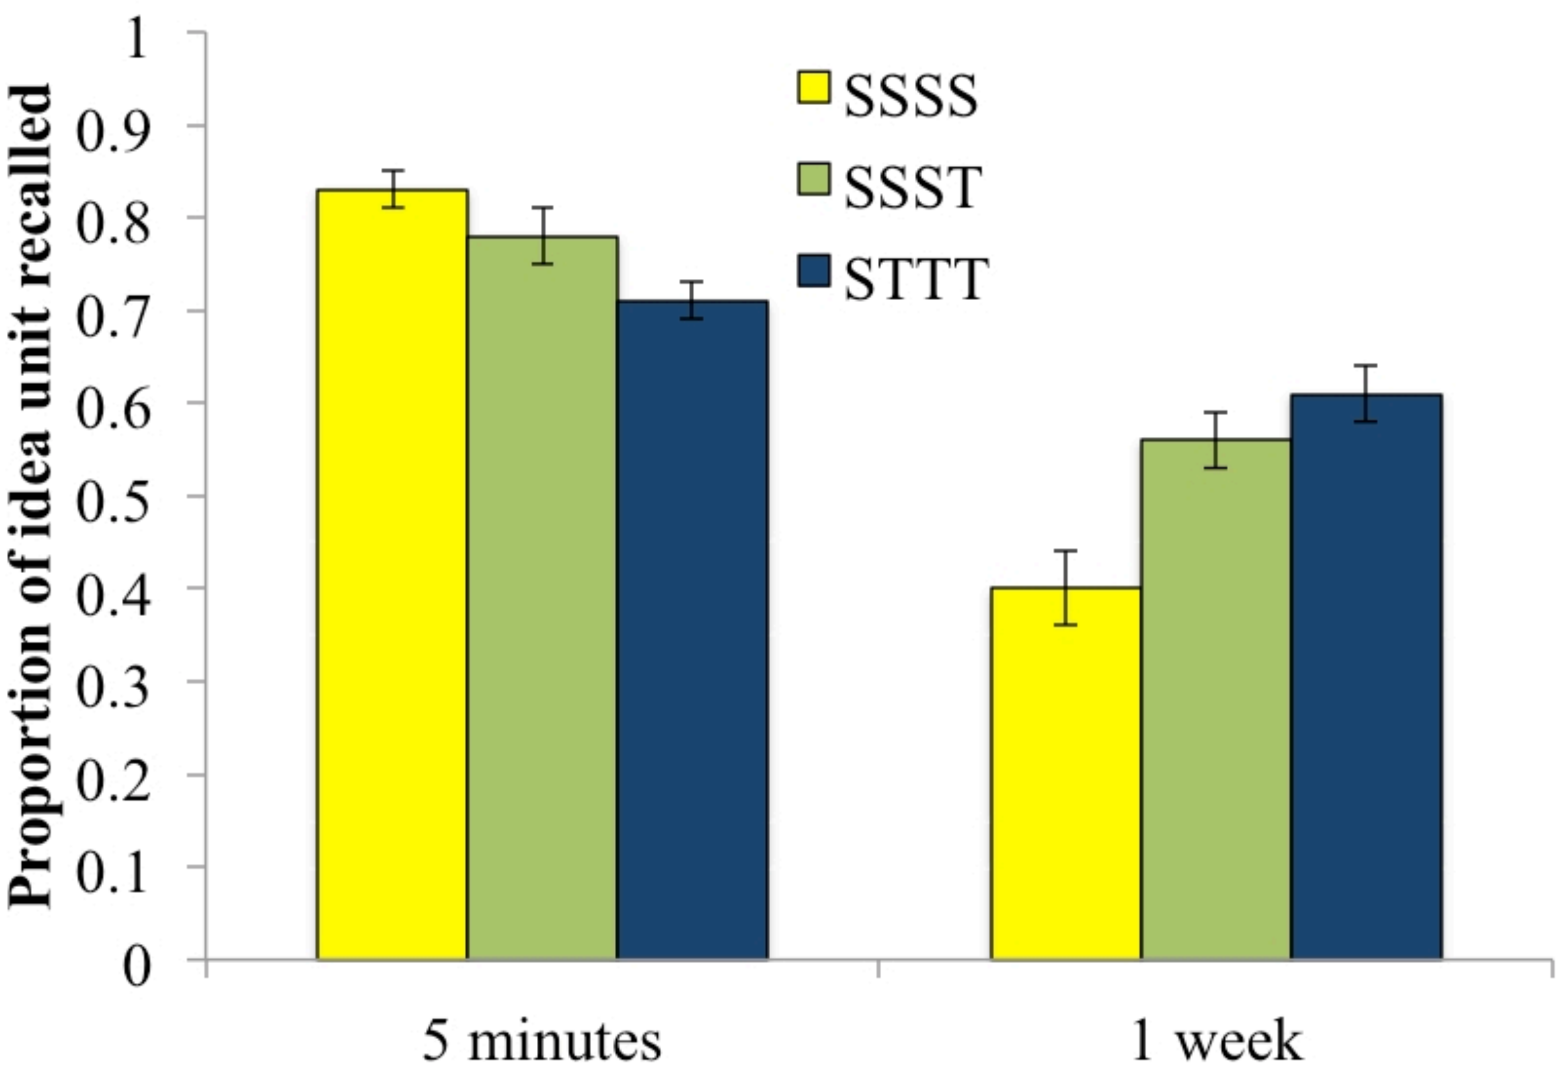

***Figure 3**. The results of the criterion free recall test in Roediger & Karpicke (2006), experiment 2. Mean proportion of idea units from the prose passage (and standard errors of the mean) show that while those who studied more recalled more at a short delay, those who took more tests retained more of the information after a longer delay.*

The testing effect is incredibly robust and has been demonstrated across a variety of stimuli and types of tests, both in cases where feedback is and is not provided (although, in cases where accuracy is very low, the benefit of testing without feedback can be absent or replaced by a benefit of rereading; for a review, see Roediger & Karpicke, 2006). One reason testing is thought to benefit learning is that the active act of retrieving information is a "memory modifier" (Bjork, 1975), strengthening what we have retrieved and weakening access to information that is in competition with the retrieved information (see Anderson, Bjork, & Bjork, 1994). If what is retrieved becomes strengthened, however, maybe it follows that testing is bad when people get the answers wrong? Maybe the wrong answer then becomes strengthened and persists, and/or interferes with subsequent learning of the correct answer (e.g., Roediger & Marsh, 2005). Concern over the negative effects of testing led, in fact, to a movement known as "errorless learning" (Skinner, 1958; Terrace, 1963).

Such fears are legitimate, but recent research has demonstrated that the benefits of testing go beyond simply strengthening what is retrieved. A good example is a study by Richland, Kornell, and Kao (2009). Across five experiments, participants were asked to study an essay about vision. In the experimental, pretest condition, participants were tested on concepts embedded in the passage, before reading the passage. Their posttest performance was significantly better than that of participants in an extended study condition, who did not answer test questions before reading, but instead used that time for extra study. This benefit was maintained even when analyses was restricted to only those items that the pretested group had answered incorrectly prior to study. In other words, taking time out of study to generate wrong, competing answers enhanced future learning. Richland *et al.* further eliminated the possibility that those who took the pretest were better able to attend to the important information by emphasizing the tested concepts through the use of bolded and italicized keywords in the extended study condition and allowing the extended study participants to read (but not answer) the same pretest questions that the experimental condition answered.

One interpretation of these findings is that activating the semantic network associated (through pretesting) with the to-be-learned topic allows for the subsequently studied information to be more elaborately encoded. Support for this view comes from research employing a simplified procedure introduced by Kornell, Hays, and Bjork (2009), one in which participants are asked to learn a list of weakly associated word pairs (such as Frog: Pond). For half of such word pairs, participants have to first try to predict the target word before being shown the intact cue-target pair; for the other pairs, participants simply study the intact cue-target pair. Despite the fact that the pairs are selected so that participants' predictions are virtually always wrong, and despite the fact that the time available to study a given intact pair is reduced by virtue of having to first try to guess the upcoming to-be-remembered target word, later performance on a cued-recall test is reliably better when participants try to predict the to-be-learned response.

It is crucial, though, and consistent with the semantic-activation idea, that the to-be-learned response bears a semantic relationship to the cue. Huelser and Metcalfe (2012) found that when the pairs are unrelated (e.g., Frog-Bench, where any semantic activation generated by "frog" would be unhelpful for encoding "bench") the benefit of trying to predict the response disappears. Importantly, Huelser and Metcalfe (2012) also asked participants at the end of their experiment using related pairs whether they thought having to first predict the to-be-remembered response before seeing that response helped or hindered their remembering that response on the final test, versus simply being shown the intact pair. Even though their own performance exhibited a benefit of the prediction condition, participants judged the pure study condition to be better for learning.

The benefit of making errors has furthermore been demonstrated in the classroom. In 7th grade mathematics classrooms from three Singapore public schools, Kapur and Bielaczyc (2012) demonstrated the benefit of what they called "productive failure." Half the classes were taught in the traditional way ("directed instruction"), cycling through seven, six, or four periods in schools A, B, and C, respectively (the variation was a result of what was afforded by each school's structure) of classroom instruction, practice, homework and feedback. The other half of the classes spent the first six, four, and two periods, respectively, working in triads on complex problems without any instruction from the teachers. In this "productive failure" condition, very few of the groups (16%, 7% and 0%) ever reached the correct solutions. Thus, these students in fact spent the majority of the class time generating errors. Compare that to the homework performance of the directed instruction condition, which averaged 91-93 percent. In the "productive failure" classrooms, teachers stepped in to provide answers in only the last one or two periods and did so by first eliciting students' failed methods and then drawing attention to the critical features of each failed solution before presenting the correct solution. On the final post-intervention test, students in the productive failure condition performed significantly better than students in the directed instruction condition on complex and graphically represented questions (i.e., demonstrating greater transfer of knowledge).

It is important to note, however, that in the productive failure classrooms, teachers did not simply tell students they were incorrect and then tell them the correct answer. Rather, they spent time discussing why incorrect solutions might have been chosen and why the correct solution was more appropriate. Similarly, laboratory experiments investigating the benefits of pretesting with text passages and simple word pairs indicate that errors should be related to the to-be-learned topic in order for there to be a benefit of making errors (e.g., Huelser & Metcalfe, 2012; Kornell, 2014; but see Potts & Shanks, 2014, for an example of a pretesting benefit even when errors are unrelated). One might argue, therefore, that is it not making errors per se that benefits learning, but rather, learning is benefited by the elaborate processing (for example, pretesting helps relate new knowledge to prior knowledge) that is involved when errors (and correct responses) are generated.

In addition to the more elaborated processing that may be possible, or even contingent on, having made a prediction error, it may be that there is, in fact, something special about the errors themselves. Making errors may be very useful, for example, in helping the learners understand what conceptual mistakes they may be prone to making in the future. If a student is simply told the answer, it may appear obvious when that answer or explanation is in front of them. At a later time when that information is needed, however, prior misconceptions that were not highlighted during learning may re-emerge.

Why Do We Not Already Appreciate What Optimizes Learning?

As mentioned earlier, a basic impediment to accurate metacognition is that current performance is an unreliable index of learning. As we have illustrated, there are many cases where current performance is not only unreliable, but also misleading. Both massing (studying and re-studying the same information without breaks) and blocking (completing study of one concept before moving onto the next) practice, for example, can lead to better performance in the short-term (and make learning feel easier) as compared to spacing (taking breaks) and interleaving practice (mixing up concepts), which can optimize long-term retention and transfer.

In addition to the fact that current performance can be misleading, it is also the case that our subjective experiences (especially that of *ease*) as a learner can be misleading. The sense of perceptual fluency we gain when reading something multiple times, for example, can be mistaken for understanding and assumed to be a reliable measure for how recallable the information will be later. Similarly, how readily information is recalled at one point in time, in the presence of certain cues, can be interpreted as a measure of how recallable that information will be at a later time, in the presence of different cues.

In addition to the fact that current performance can be misleading, it is also the case that our subjective experiences as a learner can be misleading. The sense of perceptual fluency we gain with reading something multiple times, for example, can be mistaken for understanding and a measure for how recallable the information will be later, and how readily information is recalled at one point in time, in the presence of certain cues, can be taken as a measure of how recallable that information will be at a later time, in the presence of different cues. Also, taking tests (as compared to simply studying material), particularly when one makes mistakes, can make it feel as though little learning is occurring.

There are good reasons to use indices such as perceptual fluency or retrieval fluency as guides for learning—because fluency is often an indicator of learning. If something feels easy, it may be because we really do know it! Thus, fluency can be a useful heuristic. Issues arise however, when the experience of fluency is caused by factors other than learning. Rereading the same passage 10 times, for example, will definitely lead to feeling of perceptual fluency, but the gain in learning and understanding may be minimal. Fluency, when it arises from factors unrelated to learning, can lead to illusions of competence, which then have cascading effects on students' monitoring, control, and ultimately, learning.

A good example of learners being influenced by perceptual fluency is a study by Rhodes and Castel (2008), who presented participants with a list of words, half of which were written in small font, and half of which were written in a large font. Even though actual recall did not differ between large and small font words, participants judged the large font words to be more memorable than the small font words. Reder and Ritter (1992) demonstrated that priming can also lead to the experience of fluency and increase participants' "feeling of knowing." Pre-exposing participants to certain key terms (e.g., *golf*, *par*) that subsequently appeared in a set of general-knowledge questions (e.g., *What term in golf refers to a score of one under par on a particular hole?*) led to a "feeling of knowing:" Participants were faster at estimating whether they could answer the question and more likely to judge that they could answer the question, even though such priming did not change their actual ability to answer a given question.

Retrieval fluency (how readily information "comes to mind") can be another misleading indicator of learning. In a study by Benjamin, Bjork and Schwartz (1998), for example, participants were asked a series of 20 very easy trivia questions, such as "Who was the first president of the United States?" and asked to hit the Enter button as soon as the answer came to mind. After answering each question, the participants were asked to judge the likelihood that they would be able on a later free-recall test to recall having provided that answer ("George Washington" in this example). Critically, the participants were told that they would not get the questions again but would simply be given a blank sheet of paper on which they were to write down as many of the 20 answers they provided as they could. The results were that the faster an answer came to mind, the more likely participants thought they would later be able to free recall having given that answer, whereas the actual likelihood of recalling an answer was the opposite: The longer an answer took to come to mind, the more likely they were to recall that answer on the final test. That is, the more effort that was put into generating an answer, the more likely that answer was to be later recalled in the absence of the trivia question itself. Participants, however, apparently relied on a heuristic that what is more readily recalled now will be more readily recalled in the future.

In fact, in the presence of the correct answers, one might experience an illusion of competence. Most studying takes place with the textbook and notes open, and in the presence of the answers, it can be hard to appreciate how difficult it will be to retrieve the information during the test when those notes and textbooks are not present. In the classroom, too, when students are passively listening to the lecturer, with notes in front of them and on the projector screen, they become prone to over-estimating how much they actually understand (that is, they can experience an illusion of competence). When students are asked to explain back what they learned in class, for example, they typically struggle to explain the concepts that have just been presented.

This under-appreciation of the difficulty of later being able to retrieve an answer that is present now, but will be absent and required at test, has been labeled "foresight bias" by Koriat and Bjork (2005). In one of their experiments, for example, Koriat and Bjork (2005) asked participants to learn a number of cue-target word pairs so that, later, when presented the cue word, they would be able to retrieve the target word. Some of the pairs had a very strong forward association, from cue to target, but only a weak backward association (e.g., *lamp-light*), whereas the opposite was true for other pairs. The participants were asked to judge, pair by pair, the likelihood that they would be able, on a later test, to recall the target word, given the cue word. Actual recall was significantly higher for forward pairs, because in the case of backwards pairs, such as *light-lamp*, the cue on the final test (*light-?*) triggers many other possible responses, such as *dark* or *heavy*. The participants, however, made much the same prediction for forward and backward pairs. That is, they suffered from a foresight bias—the inability to think ahead in time to the point where the correct response would be absent and in competition with other words associated to the cue.

In subsequent research, Koriat and Bjork (2006) explored ways in which students might be taught to avoid foresight bias. They found that giving learners the experience that backward-associated pairs are more difficult than forward-associated pairs enabled them to make more accurate judgments about those specific pairs; only the combination of experience- and theory-based (receiving an explanation of the asymmetric relationship between cue and target words) de-biasing techniques transferred learning to new word pairs.

The "curse of knowledge."

Finally, instructors themselves are also prone to problems of fluency, not just in judging students' learning from performance, but also in their own teaching. Instructors have to be mindful of the difference between what they know and what their students know. One might think that an expert should really make the best instructor: The expert knows the material inside out, has a strong conceptual understanding of the information, and therefore, should be able explain ideas and concepts in the most eloquent way. What is obvious to an expert, however, is not necessarily going to be obvious to a learner. The expert may not understand the misconceptions and barriers that novices face. Indeed, Piaget (1962) remarked that 'Every beginning instructor discovers sooner or later that his first lectures were incomprehensible because he was talking to himself, so to say, mindful only of his point of view. He realizes only gradually and with difficulty that it is not easy to place one's self in the shoes of students who do not yet know about the subject matter of the course' (p. 5).

One might try Newton's (1990) study as a thought experiment to illustrate the "curse of knowledge." First, tap out the rhythm of a well-known song (e.g. "London Bridge") to a listener. How likely do you think that the listener would correctly identify the song? Newton (1990) found that while the tappers (who chose from a list of 25 well-known tunes) estimated that roughly half of the listeners would successfully identify the song, the reality was that only 2.5 percent of the listeners were able to do so. What was obvious to the tappers—who could hear the tune in their heads very clearly (and perhaps even hear the lyrics and full orchestration)—was not at all apparent to the listeners who heard only a series of atonal and irregular taps. As teachers, it is crucial to understand that we are the tappers and our students are the listeners.

Implications for Instructors and for Learners

What then are the implications of cognitive psychology research for instructors and for learners? There is a certain societal assumption, particularly in college classrooms, that an instructor transfers information to students who then store that information, in ways akin to a recording device. One of the most important principles of effective learning, however—perhaps the single most important principle—is that learners must be active participants in the learning process.

If information is to be well learned, spacing out study—for example, returning to old material every so often—is necessary. Spacing engages retrieval processes, and retrieval is a powerful memory modifier—the act of retrieval strengthens what we retrieve. By the same token, low-stakes or no-stakes tests should be frequently used as pedagogical tools, rather than simply for assessment. Testing is something that both instructors and learners can implement: Instructors can introduce low-stakes tests in their courses, while learners can engage in self-testing. There are a myriad of ways in which learners can self-test: Attempting practice tests offered by the instructor, using flashcards (to test themselves, not to simply read), using end-of-chapter questions, creating their own questions, posing questions in study groups, and so on.

Another technique that is especially beneficial for inductive learning is intermixing study or practice of related concepts (i.e., interleaving, rather than blocking study). At the end of chapter, learners might practice questions both related to that chapter as well as related to preceding chapters. Interleaving practice enhances discrimination between related concepts. Additionally, especially when we consider that in exams and in real life, students need to be able to know when to apply which concepts (without the aid of having just read the relevant textbook chapter), the value of discrimination becomes clear.

In addition to understanding specific learning techniques they should use, learners should also be cognizant of the discrepancy between current performance and long-term retention and mindful of their interpretation of fluency. For example, learners should be aware that getting all the practice questions correct immediately after studying a chapter might not be an accurate measure of learning. They should, instead, wait for a period of time, perhaps study some other information, and then attempt those practice questions (without consulting the textbook or notes). This strategy could be a way of avoiding misattributions of fluency and illusions of competence, as well as incorporating beneficial learning strategies.

Efficient Learning Is Not Easy Learning: Challenging Counterproductive Assumptions

Beyond the evidence that people have faulty mental models about how they learn, and thus do not know how best to manage their own learning, there are also some societal assumptions that can be barriers to effective learning. One such assumption, which we have already mentioned, is that an instructor's responsibility is to transmit information and a student's responsibility is to record that information: That is, the instructor talks and students are expected to remain silent and absorb knowledge. As we have emphasized, however, one key to effective learn is that students must be active participants in their own learning.

The role of errors.

Another counter-productive assumption has to do with the meaning and role of errors. Errors play a critical role in our learning, yet errors and mistakes, rather than being viewed as an important component of effective learning, are often assume to reflect inadequacies of the instructor, the student, or both. The assumption that errors are to be avoided is related to another counter-productive assumption: that learning should be easy. Testing, spacing and interleaving (as opposed to rereading, massing and blocking, respectively) do not, however, make learning easy. In fact, they make learning more effortful by virtue of requiring more active processing. It is important to note here, however, that not all forms of effort are productive and not all difficulties are "desirable." Difficulties that engage deeper processing are good; difficulties that simply allow learners to wallow in frustration are not.

Over-attributing differences in performance to innate abilities.

In general, in our view, differences in performance across individuals tend to be over-attributed to differences in innate abilities and under-attributed to differences in experience, effort, and practice. When we are struggling to learn it is easy to turn blame on our own aptitude and innate abilities, or turn the blame on to an instructor's teaching style. In fact, the "styles of learning" idea—that one is a visual learner and learns best visually, for example—is very appealing: "If only teaching was presented just the right way for me, then I would learn." There is little evidence, however, to support the existence of individual learning styles (Pashler, McDaniel, Rohrer, & Bjork, 2009), and, in our view, the styles idea can be counterproductive with respect to learning. Individual differences are, of course, important: The knowledge and assumptions we bring to new learning, as well as our motivations, aspiration and expectations of learning, do matter—and matter greatly. What is critical to appreciate, however, is what we all share: a remarkable capacity to learn.

Concluding Comment

The most important message is that learners should break away from the misconception that the most effective ways of learning are those that make learning easy. The experience of having to expend effort, generate errors, or work hard to achieve understanding should not be interpreted as evidence of one's inadequacy as a learning, but, instead, as important steps towards actual long-term learning and comprehension. In short, the good news is that there exists a great potential to upgrade self-regulated learning. Far from the notion of individual "styles of learning" that pose challenges of requiring individually tailored learning methods, there are learning techniques and principles that cognitive psychology has demonstrated can benefit all of us.

References

Anderson, M. C., Bjork, R. A., & Bjork, E. L. (1994). Remembering can cause forgetting: Retrieval dynamics in long-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*, 1063-1087.

Appleton-Knapp, S. L., Bjork, R. A., & Wickens, T. D. (2005). Examining the Spacing Effect in Advertising: Encoding Variability, Retrieval Processes, and Their Interaction. *Journal of Consumer Research, 32*, 266-276.

Bahrick, H. P., Bahrick, L. E., Bahrick, A. S., & Bahrick, P. E. (1993). Maintenance of Foreign Language Vocabulary and the Spacing Effect. *Psychological Science, 4*, 316-321.

Balota, D. A., Duchek, J. M., & Paullin, R. (1989). Age-related differences in the impact of spacing, lag and retention interval. *Psychology and Aging, 4*, 3-9.

Battig, W. F. (1972). Intratask interference as a source of facilitation in transfer and retention. In J. F. Voss (Ed.), *Topics in learning and performance* (pp. 131–159). New York: Academic Press.

Benjamin, A. S., Bjork, R. A., & Schwartz, B. L. (1998). The mismeasure of memory: When retrieval fluency is misleading as a metacognitive index. *Journal of Experimental Psychology: General, 127*, 55-69.

Benjamin, A. S., & Ross, B. H. (2011). The causes and consequences of reminding. In A. S. Benjamin (Ed.), *Successful remembering and successful forgetting: A Festschrift in honor of Robert A. Bjork* (pp. 71-88). New York, NY: Psychology Press.

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

Bjork, R. A., & Allen, T. W. (1970). The spacing effect: Consolidation or differential encoding? *Journal of Verbal Learning and Verbal Behavior, 9*, 567-572.

Bjork, R. A. (1975). Retrieval as a memory modifier. In R. Solso (Ed.), *Information processing and cognition: The Loyola Symposium* (pp. 123-144). Hillsdale, NJ: Lawrence Erlbaum Associates.

Bjork, R. A. (1994). Memory and metamemory considerations in the training of human beings. In J. Metcalfe and A. Shimamura (Eds.), *Metacognition: Knowing about knowing* (pp. 185-205). Cambridge, MA: MIT Press.

Bjork, R. A. (2011). On the symbiosis of learning, remembering, and forgetting. In A. S. Benjamin (Ed.),*Successful remembering and successful forgetting: a Festschrift in honor of Robert A. Bjork* (pp. 1-22). London, UK: Psychology Press.

Bjork, R. A., Dunlosky, J., & Kornell, N. (2013). Self-regulated learning: Beliefs, techniques, and illusions. *Annual Review of Psychology, 64*, 417-444.

Carvalho, P. F. & Goldstone, R. L. (2014). Putting category learning in order: Category structure and temporal arrangement affect the benefit of interleaved over blocked study. *Memory & Cognition, 42*, 481-495.

Cepeda, N. J., Pasher, H., Vul, E., Wixted, J. T., & Rohrer, D. (2006). Distributed practice in verbal recall tasks: A review and quantitative synthesis. *Psychology Bulletin, 132*, 354-380.

Cepeda, N. J., Vul, E., Rohrer, D., Wixted, J. T., & Pashler, H. (2008). Spacing effects in learning: a temporal ridgeline of optimal retention. *Psychological Science, 19*, 1095-1102.

Cohen, M. S., Yan, V. X., Halamish, V. & Bjork, R. A. (2013). Do students think that difficult or valuable materials should be restudied sooner, rather than later? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 39*, 1682-96.

Ebbinghaus, H. E. (1885/1964). *Memory: A contribution to experimental psychology (Henry A. Ruger and Clara E. Bussenius, Trans.)*. New York: Dover.

Estes, W. K. (1950). Statistical theory of distributional phenomena in learning. *Psychological Review, 62*, 369-77.

Hartwig, M. K., & Dunlosky, J. (2012). Study strategies of college students: Are self-testing and scheduling related to achievement? *Psychonomic Bulletin & Review, 19*, 126-134.

Huelser, B. J., & Metcalfe, J. (2012). Making related errors facilitates learning, but learners do not know it. *Memory & Cognition, 40*, 514-527.

Kang, S. H. K., & Pashler, H. (2012). Learning painting styles: Spacing is advantageous when it promotes discriminative contrast. *Applied Cognitive Psychology, 26*, 97-103.

Kapur, M., & Bielaczyc, K. (2012). Designing for productive failure. *Journal of the Learning Sciences, 21*, 45-83.

Koriat, A., & Bjork, R. A. (2005). Illusions of competence in monitoring one's knowledge during study. *Journal of Experimental Psychology: Learning, Memory and Cognition, 31*, 187-194.

Koriat, A., & Bjork, R. A. (2006). Mending metacognitive illusions: A comparison of mnemonic-based and theory-based procedures. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32*, 1133-1145.

Kornell, N. (2014). Attempting to answer a meaningful question enhances subsequent learning even when feedback is delayed. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*, 106-14.

Kornell, N., & Bjork, R. A. (2007). The promise and perils of self-regulated study. *Psychonomic Bulletin & Review, 14*, 219–224.

Kornell, N., & Bjork, R. A. (2008). Learning, concepts, and categories: Is spacing the "enemy of induction?" *Psychological Science, 19*, 585–592.

Lee, T. D. (2012). Contextual interference: generalizability and limitations. In N. J. Hodges and A. M. Williams (Eds.), *Skill acquisition in sport: Research, theory, and practice II (pp. 77-93).* London: Routledge

Kornell, N., Hays, M. J., & Bjork, R. A. (2009). Unsuccessful retrieval attempts enhance subsequent learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *35*, 989-998.

Newton, L. (1990). *Overconfidence in the communication of intent: Heard and unheard melodies*. Unpublished doctoral dissertation, Stanford University, Stanford, CA.

Pashler, H., McDaniel, M., Rohrer, D., & Bjork, R. A. (2009). Learning Styles: Concepts and Evidence. *Psychology Science in the Public Interest, 3*, 105-119.

Peterson, L. R., & Peterson, M. J. (1959). Short-term retention of individual verbal items. *Journal of Experimental Psychology, 58*, 193-198.

Piaget, J. (1962). Commentary. In L. Vygotsky, *Thought and Language*. Cambridge, MA: MIT Press.

Potts, R., & Shanks, D. R. (2014). The benefit of generating errors during learning. *Journal of Experimental Psychology: General, 143*, 644-67. doi:10.1037/a0033194

Reder, L. M., & Ritter, F. (1992). What determines initial feeling of knowing? Familiarity with question terms, not with the answer. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18*, 435-451.

Rhodes, M. & Castel, A. (2008). Memory Predictions Are Influenced by Perceptual Information: Evidence for Metacognitive Illusions. *Journal of Experimental Psychology: General, 137*, 615–625. doi:10.1037/a0013684.

Richland, L., Kornell, N., & Kao, L. (2009). *Journal of Experimental Psychology: Applied, Vol 15*(3), 243-257. doi: 10.1037/a0016496.

Roediger, H. L., III & Marsh, E. J. (2005). The positive and negative consequences of multiple-choice testing. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *31*, 1155-1159.

Roediger, H. L., III & Karpicke, J. D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science*, *17*, 249–255.

Rohrer, D., & Taylor, K. (2007). The shuffling of mathematics problems improves learning. *Instructional Science*, *35*, 481–498.

Shea C. H., Lai, Q., Black, C., Park, J. (2000). Spacing practice sessions across days benefits the learning of motor skills. *Human Movement Science, 19*, 737-760.

Shea, J. B., & Morgan, R. L. (1979). Contextual interference effects on the acquisition, retention, and transfer of a motor skill. *Journal of Experimental Psychology: Human Learning and Memory, 5*, 179–187.

Simon, D. A., & Bjork, R. A. (2001). Metacognition in motor learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*, 907-912.

Skinner, B. F. (1958). Teaching machines. *Science, 128*, 969-977.

Smith, S. M., Glenberg, A. M., & Bjork, R. A. (1978). Environmental context and human memory. *Memory & Cognition*, *6*, 342-353.

Tauber, S. K., Dunlosky, J., Rawson, K. A., Wahlheim, C. N., & Jacoby, L. L. (2013). Self-regulated learning of a natural category: Do people interleave or block exemplars during study? *Psychonomic Bulletin & Review, 20*, 356-63

Terrace, H. S. (1963). Discrimination learning with and without "errors." *Journal of the Experimental Analysis of Behavior, 6*, 1–27.

Thios, S. J., & D'Agostino, P. R. (1976). *Effects of repetition as a function of study-phase retrieval. Journal of Verbal Learning and Verbal Behavior, 15*, 529-536.

Zechmeister, E. B., & Shaughnessy, J. J. (1980). When you know that you know and when you think that you know but you don't. *Bulletin of the Psychonomic Society, 15*, 41-44.