Unlearning to Learn: What the Science of Learning Teaches Us About Pedagogy

An Excerpt from *Building the Intentional University*

MINERVA®

Contents

We at Minerva are pressing the reset button on higher education. We took a step back and asked, **given what we know about the science of learning, what students need to know to succeed in the twenty-first century, and the capabilities of modern technology, what should we teach and how should we teach it?**

A key ingredient to being comfortable with the Minerva pedagogy is being open to new experiences. Part of this openness involves being willing to examine assumptions and habits and being willing to unlearn old ways of encountering something new and better. The force of this observation has become evident as we've developed a new kind of curriculum, guided by empirical findings and first principles. Our curriculum is distinctive in three respects: in what we teach, how we teach it, and the means of delivery. Each of these features has required our students and faculty to face new challenges.

In this paper, which aggregates excerpts from our book *Building the Intentional University (MIT Press, 2017)*, we focus on the second element, namely, how we teach. We explore the tenets of the science of learning and how it informs our pedagogical approach, which we call Fully Active Learning.

The Science of Learning

The science of learning encompasses findings in a wide range of areas, including discoveries about how humans perceive, organize, and store information and then subsequently retrieve that information from memory. We've learned a tremendous amount about how humans process and store information, and that knowledge can be used systematically in education to help students master the material they are taught.

Oddly, although the science of learning matured decades ago, it is rarely used to facilitate teaching. Instead, lectures are still a common way to teach, but we need to distinguish between teaching and learning. Teaching focuses on information transmission; learning is about knowledge acquisition. On the face of things, the two activities should be completely aligned, but typically they are not. Teaching is often done in a way that is convenient and efficient for the professor, with little thought given to how best to facilitate student learning. Lectures are a superb way to teach: a single instructor can lecture to ten thousand people as easily as to ten. But study after study has documented that lectures are a terrible way for students to acquire information.

Why, then, are lectures still the dominant mode of teaching in most universities? Part of the problem may be that those doing the teaching do not understand enough about the science of learning to take advantage of it. This is not to suggest that other problems are not also prevalent, such as the economics of universities, incentive structures, and institutional rigidity, but certainly most faculty who care about being effective instructors would benefit from a more thorough understanding of the science of learning.

Oddly, although the science of learning matured decades ago, it is rarely used to facilitate teaching.

Maxims

We summarize two overarching principles (called "maxims") and then consider sixteen specific principles that fall under them. Different reviewers have organized the literature differently, producing different numbers of principles. For example, Graesser, Halpern, and Hakel (2008) identify twenty-five principles, whereas Willingham (2009) identifies only nine. The differences appear to arise primarily from what principles are considered special cases or variants of other principles.

Maxim I: "Think it through"

The first maxim is "Think it through." The key idea is very simple: the more you think something through, paying attention to what you are doing, the more likely you are later to remember it.

This maxim is at the core of your ability to recall facts and figures from a newspaper article you read, even though you didn't try to memorize them. You stored the material in memory simply because you paid attention and thought it through. Incidental learning is learning that occurs without consciously trying to acquire the knowledge; it occurs as a by-product of the cognitive processing that is used to understand, analyze, or synthesize.

Maxim II: "Make and use associations"

The second maxim is "Make and use associations." Associations not only help us organize material so that it is easy to store in memory, they also give us the hooks that will allow us later to dig the material out of memory, to recall it.

A dramatic demonstration of the power of using associations to organize material was reported by Ericsson, Chase, and Faloon (1980). They asked an undergraduate student to commit to coming into the lab at least three times per week, and he did so for about a year and a half. At each session, the researchers simply read him a sequence of random digits, one digit per second, and asked him to repeat them back. They started with a single digit, which he correctly recalled. They then gave him two other randomly selected digits, which he recalled, and then three, and so on, increasing the length of each new list until he failed to recall

PRINCIPLES

Maxim I: “Think it through”

Evoking deep processing

Using desirable difficulty

Eliciting the generation effect

Engaging in deliberate practice

Using interleaving

Inducing dual coding

Evoking emotion

the entire sequence (eight digits, on that first day). Each session began where the previous one had left off, with a new list of that length (with a new combination of random digits). Every set consisted of a new set of digits; he wasn’t given practice learning the same set over and over. When the study finally ended, this participant could recall a list of seventy-nine random digits!

How did he do this? As it happened, the participant in the study was a long-distance runner who had run numerous marathons. He associated the random digits with times for particular segments of races.

Principles that underlie “Think it through”

1

Evoking deep processing.
The more cognitive operations one performs while paying attention to such operations, the more likely it is that one will later recall that information (Craig et al., 2006; Craik & Lockhart, 1972). This is the most obvious implication of the maxim “Think it through.” For example, if you formulate an example of how every one of these principles can be used in a specific situation, you will remember them much better than if you simply read and understand them.

2

Using desirable difficulty.
We can think of this as the Goldilocks rule (not too hot, not too cold—just right!). Learning is best when the task is not so easy as to be boring but not so hard as to be over the learner's head (Bjork, 1988, 1999; VanLehn et al., 2007). To get the most out of thinking it through, the person needs to be as engaged as possible—no more, no less.

3

Eliciting the generation effect.
Simply recalling information—especially when effort is required—strengthens memory for that piece of information; the mere act of digging information out of memory reconstructs and strengthens the mental representation of the information. For example, a consequence of this principle is that frequent testing can enhance learning if it leads learners to recall relevant information (Butler & Roediger, 2007; Roediger & Karpicke, 2006).

4

Engaging in deliberate practice.
In some cases, to learn effectively you need to pay attention to and think through specific aspects of what you are learning. In particular, feedback helps you to correct aspects of a mental representation when it isn't optimal (Brown, Roediger, & McDaniel, 2014; Ericsson, Krampe, & Tesch-Romer, 1993). Such feedback is most effective when learners use "deliberate practice." Deliberate practice occurs when you pay careful attention to mistakes and use the ways that an error differs from the correct performance to correct subsequent performance.

Frequent testing can enhance learning if it leads learners to recall relevant information.

5

Using interleaving.
Instead of just focusing on one type of problem (e.g., when doing math), it's best to intermix different types of problems. The same principle implies (but may not yet been investigated) that when learning French, it's best to do a bit of studying French, then some history, then some math, and then back to French.

This makes sense because it's easier to pay attention to something new than to sustain paying attention to the same material, extended over time. For example, all else being equal, you would probably learn the material in this chapter more effectively if you did something else after you finish this section, and returned to the second set of principles later.

6

Inducing dual coding.
If you are given a short paragraph to remember, you will recall it better if it includes some relevant illustrations. In general, presenting both verbal and visual material enhances memory. In this case, the brain stores multiple representations in memory (some verbal, some visual—which are stored in different parts of the brain), which gives you multiple shots at later digging the information out of storage (Kosslyn, 1994; Mayer, 2001; Moreno & Valdez, 2005).

7

Evoking emotion.
Leading someone to feel emotion when experiencing an event generally will enable him or her to recall that event more effectively. Emotion focuses attention and also causes the brain to devote extra resources to storing the information. Negative emotions in particular narrow attention and focus one on details.

Leading someone to feel emotion when experiencing an event generally will enable him or her to recall that event more effectively.

PRINCIPLES

Maxim II: “Make and use associations”

- Promoting chunking
- Building on prior associations
- Presenting foundational material first
- Exploiting appropriate examples
- Relying on principles, not rote
- Creating associative chaining
- Using spaced practice
- Establishing different contexts
- Avoiding interference

Principles that underlie “Make and use associations”

1

Promoting chunking.
As we saw in the case of the marathon runner who could memorize staggering numbers of randomly selected digits (Ericsson et al., 1980), you can use associations you already have in your memory to organize material into relatively few chunks (organized units). People can easily store in memory three or four chunks—and remarkably, each of these units itself can contain three or four chunks.

2

Building on prior associations.
When learning something new, the more associations you can find with information already stored in memory, the better (e.g., Bransford, Brown, & Cocking, 2000; Glenberg & Robertson, 1999; Mayer, 2001).

The fact that prior associations can be used to learn new information resolves an old conundrum: At one time researchers worried about a “paradox of the expert,” which hinged on the fact that the more you know, the easier it is to learn even more (Reder & Anderson, 1980; Smith, Adams, & Schorr, 1978). The intuition was that the more you know, the “fuller” memory should be—and hence it should be harder, not easier, to store new information. However, researchers have learned that the more information you already know, the more existing associations you can use to store new information. The more branches you have, the more leaves and fruit can be hung on this structure. Hence there's no actual paradox.

The more information you already know, the more existing associations you can use to store new information. The more branches you have, the more leaves and fruit can be hung on this structure.

3

Presenting foundational material first.
When complex information is to be acquired, learning is enhanced when a teacher takes advantage of existing associations to provide the most basic material first, and then integrates new material a bit at a time (Bransford, Brown, & Cocking, 2000; Wandersee, Mintzes, & Novak, 1994). Presenting foundational material first provides a backbone to which one can attach additional information, allowing an organized mental structure to be built up over time. For example, my presenting the two general maxims first should have given you a structure for understanding the specific principles.

4

Exploiting appropriate examples.
Abstract ideas cannot be fully understood without examples. But examples must be memorable, in part by being associated with prior information. Multiple examples of the same material must be associated with each other so that they form a cluster that is associated with the to-be-learned material. For example, when teaching the concept of far transfer, it's not enough to provide the example that debate techniques learned in class should then be used months later when arguing politics with friends. You would need a few different examples of far transfer, even though on the surface they may appear very different (Hakel & Halpern, 2005).

5

Relying on principles, not rote.
Learning typically requires not just becoming familiar with examples but also understanding the underlying principles that organize and integrate examples (Kozma & Russell, 1997; Bransford, Brown, & Cocking, 2000). For example, the key to far transfer is to distinguish between surface characteristics

(the particular example) and underlying deep characteristics (which tell you which knowledge should be transferred to the present case). For instance, the principles of debate can also be used in teaching, but that doesn't require becoming confrontational (a surface characteristic of debate) but rather being sensitive to the other person's goals and perspectives (a deep characteristic). The principles must be associated with the examples.

6

Creating associative chaining (a.k.a using story telling).
Stories are built on a series of interlocking causes and effects—this is the essence of a plot. Creating an interlocking sequence of associations that has a narrative arc—that is, a story—to integrate material will not only help you create larger chunks (stories are one way to build associations to create chunks) but, more than that, you also can use each part of the story to cue the next part when you later recall the material. Such cueing can greatly facilitate later recall of the information incorporated into the story (Bower & Clark, 1969; Graesser, Olde, & Klettke, 2002).

7

Using spaced practice.
Cramming may be an efficient way to study, but it's a bad way to learn. Trying to store information in one fell swoop leaves it vulnerable to being lost. One reason for this is that if you cram, you will have only one set of retrieval cues, the associations set up the one time you stored the information. If you instead spread out studying over time, you will associate the material with lots of different cues (such as cues in the room or rooms where you study, your feelings at the time, and thoughts you have while considering the information).

8

Establishing different contexts.
Far transfer is the holy grail of learning. As noted earlier, far transfer occurs when information learned in one context (e.g., a classroom) is retrieved and applied in a very different context (e.g., to a seemingly unrelated problem in a work environment years later). Far transfer appears to be possible in large part because one has learned a group of varied examples and has a firm grasp of the principles that underlie the relevant material (Hakel & Halpern, 2005; Van Merrienboer et al., 2006). But it also depends critically on knowing when learned information is relevant. To facilitate this, one should associate the material with numerous different contexts. For example, studying in different places will enhance your ability later to use the information in different contexts.

9

Avoiding interference.
Distinctive retrieval cues are crucial in part because they can help the learner avoid interference from other information (Adams, 1967; Anderson & Neely, 1996). Psychologists have documented two types of interference. Proactive interference occurs when material you have learned previously interferes with learning new information. For example, if you learned Spanish, you might have a problem learning that "de" is pronounced "duh" in French, not "day" as it is in Spanish. Retroactive interference occurs when learning new material impairs your ability to recall previously learned material. In the language example, once you learn the French pronunciation, you might have difficulty recalling the Spanish one. Creating distinctive retrieval cues can help you avoid both types of interference (e.g., you could associate the French pronunciation with an image of a French person having difficulty understanding why a learner is having this problem, perhaps dismissively saying "duh," and a Spanish person taking a siesta in the middle of the "day").

We at Minerva decided to take advantage of the science of learning, and so designed every one of our classes to rely on application techniques: Every class is built around active learning, and every one of our active learning exercises draws on combinations of the principles just described. If these principles are respected in how material is presented and used, students will learn effectively—sometimes without even trying to learn.

Fully Active Learning

As we use the term, **Fully Active Learning requires all students to be engaged at least 75 percent of the time while in class.** That is, rather than just professors inviting students to be involved in discussions, fully active learning hinges on activities and exercises that require students to engage in the sorts of cognitive processing that engender learning—namely, those processes mentioned above.

Fully active learning relies on specific pedagogical techniques we have developed. The heart of each lesson plan is its set of activities, which build on preclass assignments. We established a set of design practices that maximize the amount of active learning in each activity. Our guiding question is, "What is everybody else doing?" That is, for each activity, we focus not just on what the current speaker or actor (e.g., someone solving an equation) is doing but also on what the rest of the class is doing: we don't want students ever to sit passively and listen to what others are saying or doing. Rather, we want all students to be as engaged as possible for as much of the time as possible.

In the service of reaching this goal, we designed two practices: The first is to be deliberate and explicit about our pedagogical technique, and the second is to include as often as possible an explicit "engagement prompt" that tells all students what they should be doing when they are not actively producing a work product (e.g., speaking, writing, or otherwise acting).

Tenets of Fully Active Learning

Attention

Promote focus in class, using emotion to excite the imagination. Foster appropriate levels of difficulty.

Interaction

Teamwork, peer instruction, and discussion incorporated regularly to enhance motivation and accountability.

Application

Applied use of concepts and skills. Deliberate practice, over time, with diversified experiential contexts.

Association

Connecting ideas and experiences. Scaffolding learning to increase in complexity over time.

Feedback

Frequent, formative appraisals, assessment at regular intervals, avoiding high-stakes exams.

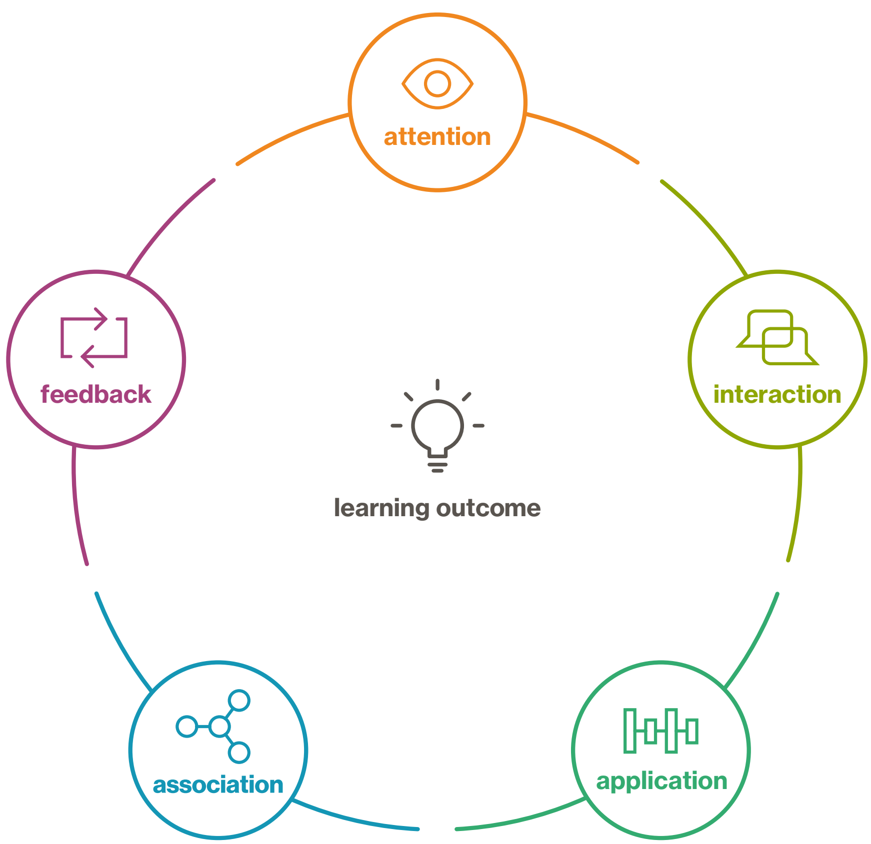

Adjusting the Way You Teach

Ample research indicates that remarkably little learning occurs on most college campuses (e.g., Arum & Roksa, 2011). Based on the principles of the science of learning, we use active learning in all classes, such as through problem solving, role-playing, engaging in debates, and the like. The specific activities are all designed to take advantage of facts about how humans learn best, which led us to design activities that require students to pay close attention and to think through problems and situations.

Adjusting to active learning

Faculty sometimes question how active learning can be better than a lecture on materials they have researched extensively and thought about for many years. In addition, faculty often have a difficult time abandoning deeply ingrained methods of teaching that they've experienced throughout their educational careers. They may not feel that they are doing a good job unless they are telling and explaining things to the students. Faculty (and students) sometimes fall prey to what we call the "illusion of learning" whereby the more notes students take, the more they are assumed to have learned. But a massive amount of research has documented that this is ineffective pedagogy: students learn best when they are actively engaged, not when they are passive recipients of a lecture.

A common way to address this problem is to have discussion sections along with the lectures. In many courses, faculty deliver two lectures and teaching assistants run a one-hour discussion section each week. But discussion is not necessarily active learning for most of the students in the class most of the time. For active learning to take place, students must actively reflect

Faculty (and students) sometimes fall prey to what we call the "illusion of learning" whereby the more notes students take, the more they are assumed to have learned.

on and try to use the information. Because most discussion leaders are not trained in implementing active learning, they may allow a few students to dominate or they may allow students to sit through the discussion passively. And when students do participate, the instructors may not appropriately challenge them, inducing them to engage with the material in new ways. Traditional discussion sections are often almost after- thoughts; they often are not structured or designed to facilitate learning but instead are merely opportunities for already engaged students to ask questions and express their opinions.

All instructors teaching sections of the same course rely on the same detailed lesson plan notes; these notes specify learning objectives, preparatory materials, assessments, assignments, and the in-class active learning exercises.

Team teaching

At Minerva, it's not just that faculty don't lecture and rely on active learning, they also work as a member of a team. All instructors teaching sections of the same course rely on the same detailed lesson plan notes; these notes specify learning objectives, preparatory materials, assessments, assignments, and the in-class active learning exercises. We adopted this technique (drawing on one used for many years at the Harvard Business School) to provide comparable learning outcomes for students in different sections of the same course and to ensure that the science of learning is being systematically operationalized in all sections. Using the same lesson plans helps us reduce variability across different sections of the same course, helping all students to have comparable experiences.

Faculty who come from traditional institutions, where they were in complete control of their classrooms and pedagogical methods, can find all of this structure challenging. This is especially true of faculty who studied or taught at institutions where there were no core curricula requiring this kind of intersectional coordination. However, our faculty members typically come to thrive in this structure. Every week all the faculty who teach sections of the same course meet and review the upcoming week's lesson plans, revising and updating them as needed. During the discussion they review how well the previous week's lessons went and note ways that the lesson plans should be revised for the following year. Faculty soon learn that working as members of a team helps them teach better and helps their colleagues teach better. Faculty report that this experience is gratifying.

Facilitating active learning requires a different set of skills than is required to lecture.

Despite a uniform lesson plan, all actual lessons unfold differently from one another because of differences in student responses and faculty reactions. There is also flex time built into every lesson plan, and faculty report "best practices" to their team every week— which in turn helps improve the lesson plans that will be used in the future.

Faculty effort and stimulation

The lack of complete control over the class content can also require faculty to put in more effort than they would in a traditional setting. One reason why faculty members like to lecture is because they know in advance what will unfold during the class period. Not only do they know what they are going to say, but after giving the lecture a few times, they often have a good sense of what questions the students will ask. Unlike in a Minerva seminar, faculty giving traditional lectures do not need to worry about handling challenging questions, students who are lost and vocal about their confusion, or class discussions that may veer from the learning objectives, nor do they need to worry about ensuring that all students are drawn into the activities.

Facilitating active learning requires a different set of skills than is required to lecture. Minerva faculty need to learn how to use the technology and manage the classroom activities within the allotted time. Almost anything can happen—and that initially makes some professors nervous. But we have also noticed that faculty soon find these seminars unusually stimulating. They are never bored, and they find that they learn more than they did when teaching using traditional methods.

Student engagement

Active learning also requires work on the part of the students. Taking notes during lectures is easy; having to interact with other students and think through a difficult problem or situation is not. Students must be prepared when they enter class and must remain engaged throughout every class. They cannot be passive, back-row observers. The key to active learning is that the more students pay attention and work through material, the deeper their learning becomes.

Conclusion

At Minerva we have learned how important it is to have an open mind, to be willing to unlearn old approaches and learn new ones—on the part of both faculty and students. The effort this requires is not for everyone; both faculty and students must be open to new experiences and convinced that it's worth giving the new approaches a try.

Acknowledgments

This paper takes excerpts from chapters 10, 11, and 12 from the book *Building the Intentional University: Minerva and the Future of Higher Education (MIT Press, 2017)*. The chapters are authored by Stephen M. Kosslyn, Joshua Fost, Robin B.Goldberg, Teri Cannon, and Rena Levitt.

Design and Layout by Laura Galvin

On the Cover: A neural network emerging from a brain represents the active learning principle of applying knowledge across new and unfamiliar contexts.

Photo by notorious v1ruS unsplash.com/@virupro10

***Building the Intentional University* can be purchased at:** https://mitpress.mit.edu/9780262536196/building-the-intentional-university/

For more on this or our partnerships, contact us at: info@minervaproject.com