

Learning Science & Literacy

Useful Background for Learning Designers

Contents

3 Introduction

- Why learning science?
- Why literacy?
- What happens in grades 4-13+?

6 What do we know about how people learn?

- What is learning science?
- What is learning?
- Working memory and long-term memory
- The goal of learning: developing expertise

11 Eight principles of learning science

- Get explicit instruction
- Activate prior knowledge
- Organize knowledge
- Manage cognitive load
- Practice deliberately
- Get feedback
- Gradually remove scaffolds
- Use metacognition

26 Learning to be literate

- What are the language and literacy skills that students develop in grades 4-13+?
- What do good readers do?
- What do good writers do?
- What do good communicators do?

33 In conclusion

34 References



Introduction

Hi Intentional Futures here. We've spent several months learning about learning alongside our partners at the Bill & Melinda Gates Foundation, and we're excited to share that learning with you.

In February 2015, we worked with members of the foundation's US Programs team to prepare for a conversation about improving student outcomes in literacy, particularly when it comes to older students who have already mastered the basics of phonics and decoding. We planned a lively discussion with an eclectic group including foundation staff, thought leaders, academics, technologists, and inventors. We knew everyone was coming to the conversation with a different knowledge base and wanted to help level the playing field, establishing a common foundation for the folks in the room. Many participants were interested in learning science and wanted to know more about fields like cognitive psychology but hadn't (yet!) had opportunities to read deeply on the topic. To help everyone get up to speed quickly we prepared a pre-read on learning science and literacy. That document has gone through several evolutions, the latest of which you're holding in your hands today.

We're sharing what we learned in this format because we think it can be useful to others – teachers, school leaders, parents, curriculum developers, product designers, and more. There's a wealth of information about learning science available to those who are interested, but it can be difficult to know where to start. This is a brief overview that highlights some core principles – we hope it will encourage you to apply learning science to the learning experiences you design (and perhaps even to your own life). It's certainly given us good food for thought!

At several points this document discusses the potential that technology has to support teaching and learning, but it isn't intended to put a stake in the ground about the role of technology in schools. People around the country are working tirelessly to help students master the skills they'll need for success in college, career, and life, and they're using many different approaches both in person and online. Our goal is to spark conversation about what we know today about how students learn – to share principles that are useful for everyone who's engaged in the hard work of designing learning experiences, no matter the format.

While engaged in this project, we operated as synthesizers, not primary researchers: as you read, you'll see that we've quoted liberally from other texts. The ones cited most frequently in the learning science section include Willingham's *Why Don't Students Like School: A Cognitive Scientist Answers Questions About How the Mind Works and What it Means for the Classroom*, Ambrose et al.'s *How Learning Works: Seven Research-Based Principles for Smart Teaching*, and Brown et al.'s *Make it Stick: The Science of Successful Learning*. In the literacy section we frequently quote McNamara's *Reading Comprehension Strategies* and Graham's *Effective Writing Instruction for All Students, Writing to Read*, and *WritingNext* publications. These are all great overviews that contain multiple references to scholarly articles as well as extensive bibliographies. We've made additional notes throughout about some of the specific sources we recommend for anyone who would like to dive deeper into these topics.

We hope you'll find the learning journey as stimulating as we have.

► Why learning science?

In the past century, films, radio, television, personal computers, laptops, smartboards, and tablets have all in turn been heralded as revolutionary technologies for learning (see Hess and Saxberg, 2013, p. 10-16). They've also failed to produce the massive learning gains that advocates promised. As researchers have sought to understand why, they've "discovered that... teachers were using computers [and other technologies] as quick add-ons to the existing instructionist classroom" (Cuban, 2001, as cited in Sawyer, 2014, p. 12). In other words, when technology was added on top of a fundamentally broken model, it wasn't enough to transform students' learning.

A scan of most products on the market today reveals that few edtech offerings are designed in ways that reflect what cognitive scientists and other researchers know about how people learn. This makes us nervous: if edtech isn't well-designed, rooted in research about how people learn, it's as unlikely to have an effect on student learning as film strips or radio broadcasts.

We'd like to change that. We believe in the potential technology has to improve access to effective instruction, and we would like to see it become an effective tool that's easy for educators to use in the context of a well-designed school model. Intelligent tutoring systems, adaptive software, and other tools designed in a way that takes learning science into account are starting to show significant learning gains and we'd like to build on that momentum.

If learning science can inform the design of instructional technology in ways that drive better outcomes, the opportunity is large – the majority of today's students are exposed to this kind of software in some way, shape, or form:

- In K-12, "70 percent of US teachers report that their students use educational technology to learn or practice basic skills during class at least sometimes" (Gray et al., 2009, as cited in BMGF Literacy Courseware Challenge, 2013, p. 6).

- In the higher ed sphere, "nearly four out of five US-based postsecondary students (79%) have taken at least one blended course with some online components and some face-to-face components" (EDUCAUSE, 2013 as cited in BMGF Post-Secondary Courseware Challenge, 2014, p. 5).

We want to consider ways to better leverage the potential that edtech has – this is why we'll start by acquainting the reader with the foundations of learning science.

One important note: learning science is a broad field that includes the study of cognitive processes (for example, how novices internalize new information and learn to use it). It also includes the study of motivation, mindsets, the powerful role of social factors in learning, and more. These "noncognitive" skills are essential for good learning and work in concert with the cognitive skills described in this document. For an extensive review of the literature on noncognitive skills, see the University of Chicago's *Teaching Adolescents to Become Learners: The Role of Noncognitive Factors in Shaping School Performance: A Critical Literature Review* (Farrington, 2012).

► Why literacy?

Learning science can be applied in any content area and technology can be used as a tool to support learners across many contexts. We've chosen to narrow in on ways that learning science can be used to support students as they develop skills in reading, writing, and argumentation.

Being a fluent reader, writer, and speaker is critical for success at the college level. Literacy skills such as writing are also becoming **indispensable in the workplace:**

Reports by the National Commission on Writing (2004, 2005) reveal that... the majority of both public and private employers say that writing proficiency... directly affects hiring and promotion decisions. The demand for writing proficiency

is not limited to professional jobs but extends to clerical and support positions in government, construction, manufacturing, service industries, and elsewhere.

Graham and Perin, 2007, p. 8-9

However, too many employees arrive in the workplace with skill gaps that prevent them from performing at the level their jobs demand: about 30% of government and private sector employees “require on-the-job training in basic writing skills,” and companies spend billions of dollars annually to remediate those skills (Graham and Perin, 2007, p. 9).

These skill gaps form early: 2013 National Assessment of Educational Progress (NAEP) results showed that while reading proficiency rates have climbed slowly over the last ten years, 66% of all fourth grade students—and a full 80% of lower income students—are not yet reading at a proficient level (Annie E. Casey Foundation, 2014, p. 1).

Not surprisingly, scores for twelfth graders were comparably low: 62% of students were performing below the proficient level on the 2013 exam (Nation’s Report Card, 2013). Students’ writing abilities were even further behind, with 73% of US students in grades 8 and 12 performing below a proficient level in writing (Nation’s Report Card, 2011).

This means that literacy skills are a blocker for many students. Students “who do not read proficiently by the end of third grade are four times more likely to leave school without a diploma than proficient readers” (Annie E. Casey Foundation, 2013, p. 1) and many of those who do graduate face an uphill battle in college:

An analysis of college students enrolled during the 2007-08 school year estimates that remediation needs throughout their time in college cost the nation an estimated \$5.6 billion. This figure includes \$3.6 billion in direct remedial education costs for students who did not have the skills to succeed in postsecondary course work. It also

includes an additional \$2 billion in lost lifetime wages, since students who take remedial courses are more likely to drop out of college without a degree.

Alliance for Excellent Education, 2014, p. 1

Improving students’ literacy skills is a matter of national importance, and we’re eager to apply lessons from learning science to help.

► **What happens in grades 4-13+?**

In this document, we’ll focus on the literacy skills that students need to acquire beginning around fourth grade. This is an important transition point: by fourth grade, students who are on level have typically learned to read and are now “reading to learn.” In other words, these students have mastered the basics of phonemic awareness (blending and segmenting sounds) and phonics (decoding letter combinations; being able to sound words out) and they have developed some measure of fluency (reading quickly, accurately, and with enough understanding to support comprehension).

At this stage, it becomes more important for students to develop an academic vocabulary, read grade level-appropriate texts across a variety of genres, and use their reading skills to acquire content knowledge in new subject areas. As they develop these skills, they also begin to write longer and more complex texts that inform, persuade, and entertain. Their writing is increasingly informed by—and informs—their reading. They use oral language skills to present, debate, and discuss. The skills that students begin practicing in late elementary and middle school are refined in high school and carry them into the first years of college. Students who start college in need of remediation focus on many of these same skills as part of their developmental English courses. In other words, these are skills that all students in grades 4-13+, regardless of background, need instruction and practice in order to master.

With all that said, it’s time to dive in – let’s learn about learning!

What do we know about how people learn?

► What is learning science?

Learning science refers to the interdisciplinary study of how people learn, as well as the design and implementation of effective learning environments (Nathan, 2014, p. 21). The field includes academic disciplines such as cognitive psychology, educational psychology, anthropology, linguistics, computer science, and neuroscience, as well as professional disciplines such as instructional design and educational technology. As noted earlier, this document will focus on cognitive science, which is one important subdomain within the learning sciences.

Learning scientists include both researchers and practitioners, and leaders within the learning science community are beginning to advocate a “learning engineering” mindset that would bridge research and practice, creating cycles of continuous improvement (see Hess and Saxberg, 2013).

► What is learning?

Learning results from what the student does and thinks and *only* what the student does and thinks. The teacher can advance learning only by influencing what the student does to learn.

Herbert A. Simon

Quoted in Ambrose, 2010, p. 1 (emphasis added)

Learning is “a process that leads to change, which occurs as a result of experience and increases the potential for improved performance and future learning” (adapted from Mayer, 2002 by Ambrose et al., 2010, p. 3). There are some important subtleties to unpack:

I Learning is a **process**, not a product. However, because this process takes place in the mind, we can only infer that it has occurred from students’ products or performances.

2 Learning involves **change** in knowledge, beliefs, behaviors, or attitudes. This change unfolds over time; it is not fleeting but rather has a lasting impact on how students think and act.

3 Learning is not something done to students, but rather something students themselves do. It is the direct result of how students interpret and respond to their **experiences** – conscious and unconscious, past and present.

Ambrose, 2010, p. 3

Learners must actively construct knowledge. Researchers agree that this means actively processing new information and rehearsing it in order to transfer it from working memory into long-term memory. Long-term memory is the “central, dominant structure of human cognition” (Clark, 2012, p. 9). The process of transferring knowledge from working memory (where conscious processing occurs) to long-term memory (where it can become automated and where we can draw upon it at will) is the essential process by which learning takes place, and it is key to helping novices gain expertise (Clark, 2012, p. 9).

▶ Working memory and long-term memory

The human brain stores information in both working memory and long-term memory. Our **working memory** is our conscious mind, the part of our mind where we hold the things we’re actively thinking about – there are strict limits to how much we can process in working memory at one time (more on this later). Much of what’s in our **long-term memory** is outside our conscious awareness – it’s where we store things for later use and compared with working memory, it’s almost unlimited. Many people think of long-term memory as a passive storage location for isolated bits of information. It’s much more than that – we actually rely on our long-term memory for everything from cooking to playing chess. Long-term memory contains facts (ex: the color of a polar bear), process information about how things work (ex: how plants generate energy through photosynthesis), procedures (ex: how to make a safe left turn while driving, how to multiply 18 x 7), and more (Willingham, 2009, p. 7-16). Some experts go even further: “it’s really long-term memory that defines ‘who we are,’ given that it determines what we react to immediately, instinctively, naturally, and seemingly without effort” (Hess and Saxberg, 2013, p. 41).

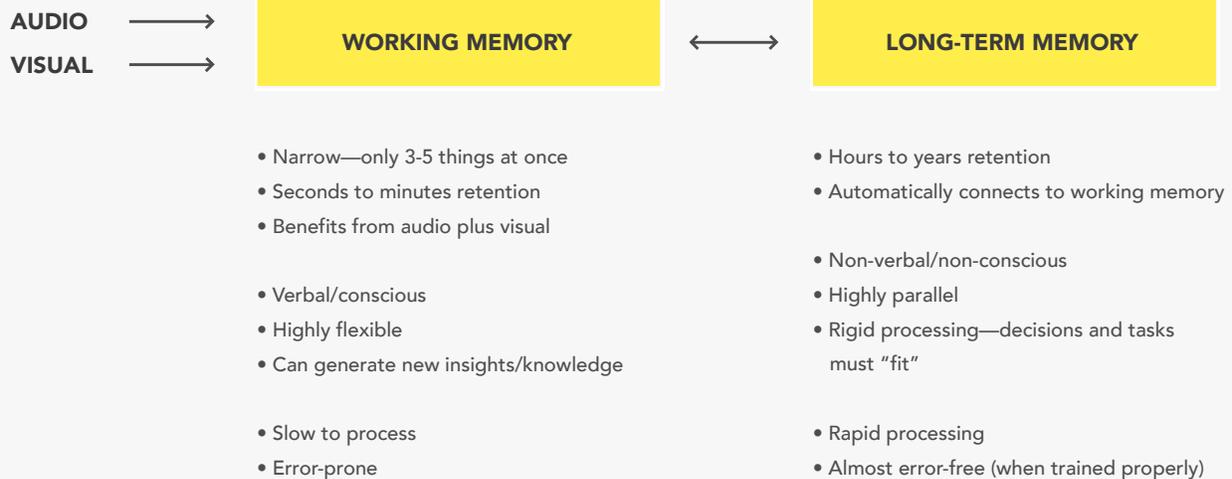


TABLE 1: A comparison of working memory and long-term memory (adapted from Hess and Saxberg, 2013, p. 40).

When we think, we pull information from our long-term memory into our working memory. **Thinking** “occurs when you combine information from the environment and long-term memory in new ways” (Willingham, 2009, p. 14).

Learning scientists believe “your memory is not a product of what you want to remember or what you try to remember; it’s a product of what you think about” (Willingham, 2009, p. 55). This means that when we process information in working memory, it’s important to think about it the way we’ll want to remember and apply it – if we need to remember the meaning or significance associated with a piece of content, then we need to spend time thinking about (processing, discussing, and reflecting on) that meaning (Willingham, 2009, p. 61).

It may seem obvious that in order for learners to get something into working memory, that they need to pay attention to it. However, paying attention is not sufficient for moving content into long-term memory. Surprisingly, having repeated exposure to content and wanting to remember it is also insufficient for building long-term memory (Willingham, 2009, p. 58-59).

It turns out that in order to transfer new knowledge and skills into long-term memory, learners have to relate the new information to what they already know, integrate it into more complex knowledge structures, and correctly rehearse new skills enough times to develop some degree of automaticity with them (see Principles #2, 3, and 5, which are introduced later in this document).

► **The goal of learning: developing expertise**

Understanding how learners progress from a novice to an expert state has been a central focus of study in the learning sciences. Becoming an expert is hard work and it takes time: learners need to believe there’s value in developing expertise and have confidence in their ability to make progress before they’re likely to put forth the effort it will take. Clearly we don’t expect young students to become experts in all the domains they study in

school – becoming an expert takes years, and most of us will become experts in very few things over the course of a lifetime. But while students may not become experts, they do need to gain mastery and develop expertise in multiple subject areas over the course of a scholarly career. In other words, one goal of a formal education is to help learners begin moving down the path from novice to expert.

Before diving into learning science-based principles about how to build mastery and acquire expertise, let’s consider what it means to be an **expert**:

There are four things experts have in common:

- 1 When it comes to many decisions and procedures in their domains, **they work quickly.**
- 2 When making routine decisions in their domain, **they rarely make mistakes** – even when the decisions are quite complex.
- 3 When they talk about what they routinely do, **they organize their conscious thinking at a completely different level** than nonexperts.
- 4 **They tackle new challenges in their domain faster and more accurately** than novices by organizing their thinking about concepts and principles differently.

Hess and Saxberg, 2013, p. 37

Experts “seem to have a sixth sense about what is important and what should be ignored” and they can be “sensitive to subtle clues that others miss” (Willingham, 2009, p. 131). They create densely connected webs of knowledge within their long-term memories – these webs are based on fundamental principles, rather than disconnected facts, and that makes it easier for them to retrieve and apply information (Ambrose et al., 2010 p. 43, 49) (see Principle #3, which will be introduced in the next section of this document).

Learning scientists distinguish between declarative knowledge and procedural knowledge. **Declarative**

The 70% rule: experts who teach can't recall about 70% of their own automated decisions and analytical strategies.

Clark, 2012, p. 544

knowledge includes facts, concepts, principles, and processes (ex: how erosion shapes the land) – this kind of factual knowledge stays relatively conscious for experts. **Procedural knowledge** is different: this is knowledge of the steps involved in completing a task. Once experts build a skill, they automate it – this means they can perform routine tasks without taxing their working memory. (Imagine an expert mathematician who had to re-explain the steps to himself every time he wanted to solve a system of equations – he'd never get anything done!)

Expert problem solvers derive their skill by drawing on the extensive experience stored in their long-term memory in the form of concepts and procedures, known as mental schemas. They retrieve memories of past procedures and solutions, and then quickly select and apply the best ones for solving problems.

Clark et al., 2012, p. 9

In other words, experts are very good at solving problems because they have automated relevant procedural knowledge in their domain. Unfortunately, this has a negative side effect when experts try to teach novices.

Learning scientists generally agree that providing explicit instruction helps novices jump start the process of developing expertise (see Principle #1 in the section that follows). However, they also agree that simply having experts explain what they know is inadequate for teaching novices – the experts in a domain aren't necessarily the best instructors for beginners. Why?

Experts don't know what they know. The fact that they have automated their procedural knowledge—which is essential for their professional practice, and is actually part of what identifies them as experts in the first place—means that they make many errors and omissions when explaining what they do to others. This phenomenon is sometimes referred to as **the 70% rule**. Researchers found that in the field of healthcare, “experts who teach can't recall about 70% of their own automated decisions and analytical strategies but must describe an approach to students and... tend to fill in their memory gaps with assumptions that are often wrong or irrelevant” (Clark, 2014, p. 544, emphasis added). They tend to focus too much on imparting declarative knowledge, ignoring critical procedural (“how to”) knowledge (Richard Clark, interview, 2014).

To overcome the 70% gap, researchers use a method called **Cognitive Task Analysis (CTA)** to understand what experts in a given domain actually do. CTAs use interviews to elicit the knowledge that experts draw upon and identify the decisions they make. By taking advantage of CTAs, learning scientists can design instruction for novices that's based on what experts actually do, rather than what they think they do.

Meta-analyses of CTA-based instruction have found performance gains up to 45%, with an average learning gain of 31% across all methods. Take one illustrative example: in a randomized controlled study of medical students in emergency medical specialist training, students who received CTA-

based training outperformed the expert-taught group by over 50% and in the subsequent year of practice, they made zero life-threatening mistakes (this compares favorably to students in the control group who made four such mistakes) (Clark, 2014, p. 543).

Discovering the 70% gap in experts' recall and using CTAs to uncover hidden knowledge, procedures, and decisions is a prime example of the value learning science can provide. By discovering gaps in what we think we know about learning and creating strategies to address those gaps, we can provide more effective instruction for the students who need it most.

Eight principles of learning science

As a quick glance at many of the sources listed in the bibliography will show, there are various sets of learning science principles afloat in the world today. Each set of principles comes as an attempt to synthesize massive amounts of learning science research, and each one highlights areas where there is consensus in the learning science community.

So why create another set? We thought it was important to use principles that focused on the processes that are key for developing language and literacy skills. We also thought it important to select a set of principles likely to yield lively discussion about the role that technology can play in creating learning experiences. We've arrived at a list of principles that takes inspiration from many of other sets we've found useful.

Each of these principles can be applied to the design of instructional technology – most edtech vendors with a product on the market today will probably be quick to tell you all about the ways they provide deliberate practice or give students just-in-time feedback. However, learning scientists agree that the devil's always in the details when it comes to instructional design: applying these principles can actually be quite complex and counterintuitive, and gaining familiarity with what cognitive science has to say about learning is likely to be helpful for everyone trying to design a learning experience.

In this document, we'll elaborate on eight key principles:

- 1 Get explicit instruction**
- 2 Activate prior knowledge**
- 3 Organize knowledge**
- 4 Manage cognitive load**
- 5 Practice deliberately**
- 6 Get feedback**
- 7 Gradually remove scaffolds**
- 8 Use metacognition**

1 Get explicit instruction

PREVIEW: Novices (but not experts!) learn the most through fully guided instructional experiences that take advantage of tools like worked examples.

All learners begin as novices, and novices learn most effectively and efficiently when they receive direct, explicit instruction (also sometimes called **fully guided instruction**). This type of instruction:

Fully explain[s] the concepts and skills that students are required to learn. Guidance can be provided through a variety of media, such as lectures, modeling, videos, computer-based presentations, and realistic demonstrations. It can also include class discussions and activities – if the teacher ensures that through the discussion or activity, the relevant information is explicitly provided and practiced.

Clark, 2012, p. 6

Fully guided instruction is carefully designed based on an analysis of the knowledge and skills students need to acquire – ideally, it’s based on an analysis of what experts do (think: CTA). This type of instruction typically includes an explicit introduction to new material, gives students an opportunity to practice (see Principle #5), provides corrective feedback (see Principle #6), and fades instructional supports as learners become more independent (see Principle #7).

Getting explicit instruction—as opposed to going through less guided instructional experiences—is particularly important for the learners who are furthest behind. In fact, in a review of approximately 70 studies that included various kinds of instruction and students with a range of abilities, researchers found that less skilled learners performed better with more guided instruction. An even more striking finding was that:

Minimally guided instruction can increase the achievement gap... A number of experiments found that less-skilled students who chose

or were assigned to less-guided instruction received significantly lower scores on post-tests than on pretest measures. For these relatively weak students, the failure to provide strong instructional support produced a measurable loss of learning.

Clark, 2012, p. 8

Worked examples are an example of a particularly effective form of explicit instruction that can be easy to apply. A worked example is “just what it sounds like: a problem that has already been solved (or ‘worked out’) for which every step is fully explained and clearly shown” (Clark, 2012, p. 9) – when going through a worked example, students do not have to discover or invent any of the steps they’ll need to take in order to succeed.

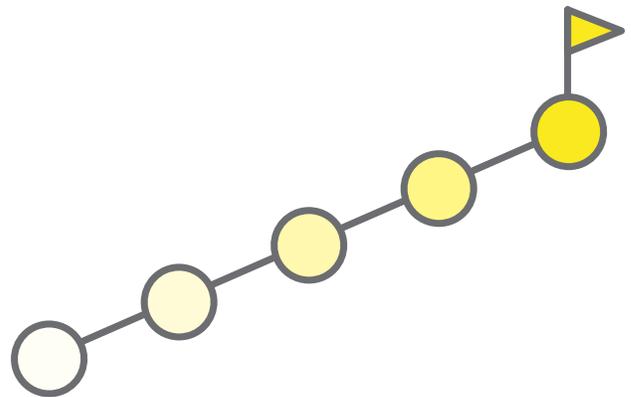


FIGURE 1: Novices need a path to follow. Unlike experts, they learn the most through fully guided instructional experiences.

Solving a problem requires searching for a solution: if “the learner has no relevant concepts or procedures in long-term memory, the only thing to do is blindly search for possible solution steps,” which “overburdens limited working memory and diverts working-memory resources away from storing information in long-term memory. As a consequence, novices can engage in problem-solving activities for extended periods and learn almost nothing” (Clark, 2012, p. 10). Direct instruction reduces the burden on working memory and directs a learner’s attention to the knowledge and skills that are to be stored in long-term memory.

One final note on explicit instruction: as we said at the outset, it's a strategy that's effective for novices. Learners fall on a spectrum from novice to expert, and as they develop greater proficiency along this spectrum, they benefit less from fully guided instruction. This is the **expertise reversal effect**: research shows that "instructional techniques [like worked examples] that are highly effective with inexperienced learners can lose their effectiveness and even have negative consequences when used with more experienced learners" (Clark 2012 p. 10). This may be because when experienced learners are provided with too much direction, they tend not to use the highly successful skills and strategies that they have already developed and automated (Richard Clark, interview, 2014).

The expertise reversal effect highlights the importance of determining which learners are truly novices, and which have developed some expertise. Early research suggests that in some domains, using a rapid pretest method may be able to identify students who are too advanced to benefit from explicit instruction, but more work needs to be done to understand how this might generalize (Kalyuga and Sweller as cited by Richard Clark, interview, 2014).

As learners master basic skills and begin to develop expertise, it's important to fade out guidance and give them opportunities to make decisions and apply their knowledge more independently. For more on this topic, see Principle #7 (scaffolding).

REVIEW: Novices (but not experts!) learn the most through fully guided instructional experiences that take advantage of tools like worked examples.

For readers who wish to know more about fully guided instruction, we highly recommend Clark, 2012.

2 Activate prior knowledge

PREVIEW: Prior knowledge is stored in chunks that can be pulled seamlessly into working memory – there, it is used as a lens for examining new ideas and as Velcro that helps them

stick. However, prior knowledge needs to be activated, sufficient, appropriate, and accurate in order to fully support a learning objective.

No student is a *tabula rasa*; rather, students come to every learning activity with knowledge they've acquired in both formal and informal settings. This knowledge "consists of an amalgam of facts, concepts, models, perceptions, beliefs, values, and attitudes" that are related to the kind of task being performed (Ambrose et al., 2010, p. 13). Prior knowledge is a **lens** through which students interpret incoming information (Ambrose et al., 2010, p. 15) as well as a kind of **Velcro** that helps new knowledge stick (Sharon Vaughn, interview, 2015). There is "widespread agreement among researchers that students must connect new knowledge to previous knowledge" in order to learn and be able to access new information in the future (Ambrose et al., 2010, p. 15).

Take reading as an example: background knowledge provides readers with context that's essential for their comprehension, particularly when they're reading from a specific content domain like science or history. Prior knowledge "provides basic vocabulary," "allows [readers] to bridge logical gaps that writers leave," and "guides interpretation of ambiguous sentences" (Willingham, 2009, p. 36). An illustration:

A clever study on this point was done with junior high school students. Half were good readers and half were poor readers, according to standard reading tests. The researchers asked the students to read a story that described half an inning of a baseball game. As they read, the students were periodically stopped and asked to show that they understood what was happening in the story by using a model of a baseball field and players. The interesting thing about this study was that some of the students knew a lot about baseball and some knew just a little. (The researchers made sure that everyone could comprehend individual actions, for example, what happened when a player got a double.) The dramatic finding... was that students' knowledge of baseball determined how much they understood of the story. Whether

they were ‘good readers’ or ‘bad readers’ didn’t matter nearly as much as what they knew.

Willingham, 2009, p. 35

Prior knowledge confers an additional benefit: it allows for **chunking**. As we’ve discussed, working memory limits the amount that we can actively think about or process at one time. But as learners store information in long-term memory, they organize it (consciously and unconsciously) into chunks: that is, they form larger units of knowledge that can be pulled seamlessly into working memory. Chunking doesn’t increase the capacity of working memory, but it does help shrink the size of the contents within a learner’s working memory. Learners “can keep more stuff in working memory if it can be chunked. The trick, however, is that chunking works only when you have applicable factual knowledge in long-term memory” (Willingham, 2009, p. 34).

To think is a transitive verb. You need something to think about.

Willingham, 2009, p. 5

Let’s use a working memory task as an example (taken from Willingham, 2009, p. 33):

1. XCN, NPH, DFB, ICI, ANC, AAX
2. X, CNN, PHD, FBI, CIA, NCAA, X

Picture a task where a typical person is asked to read the letters in one of those lists a single time, then cover them up and recall as many letters as possible. Most people perform vastly better when recalling List 2. This isn’t surprising. For the most part, List 1 appears to be a random assortment of letters – French speakers might recognize that the letters ICI form the word *ici* (“here”) and other people may have a different association with another specific letter string, but most of the content of List 1 is meaningless to most of us. Contrast this with List 2, which is full of common acronyms that most Americans have heard repeatedly in a variety of contexts – knowing these acronyms turns the random letter strings into meaningful chunks that are easier to recall and manipulate.

Readers who can chunk the individual ideas within a text into larger constructs are more likely to comprehend: imagine reading a text that introduces ideas A, B, C, D, E, and F and trying to relate those ideas to one another in order to comprehend the text’s overall meaning. That’s a lot for a reader to keep in working memory. But suppose the reader could chunk A through E into a single idea? Comprehension suddenly becomes much easier (Willingham, 2009, p. 33).

The facts and basic understandings that comprise students’ prior knowledge are the foundation on which they build more advanced skills. No educator would seriously argue that students should memorize facts and dates for their own sake, but “research from cognitive science has shown that the sorts of skills that teachers want for students—such as the ability to analyze and think critically—require extensive factual knowledge... To think is a transitive verb. You need something to think about” (Willingham, 2009, p. 25).

A few important caveats regarding prior knowledge:

- Sometimes learners' prior knowledge is **inactive**. In order for prior knowledge to be helpful, learners have to be thinking about it and pulling it from their long-term memory into their working memory – this doesn't always happen spontaneously (see Ambrose et al., 2010, p. 15-18).
- Sometimes learners' prior knowledge is accurate but **insufficient**. Learners might have holes in their declarative knowledge (knowledge of facts and concepts) or their procedural knowledge (knowledge of when and how to apply various procedures, theories, styles, or approaches). For example, "business students may be able to apply formulas to solve finance problems but not explain their logic or the principles underlying their solutions" (see Ambrose et al., 2010, p. 18-19).
- Sometimes learners' prior knowledge is **inappropriate**. Knowledge from one context (whether technical or everyday) can interfere in another context if learners apply it inappropriately. For example, students who think of writing "as a 'one-size-fits-all' skill... [might] misapply conventions and styles from their general writing classes to disciplinary contexts in which they are not appropriate" – this might mean "applying the conventions of a personal narrative or an opinion piece to writing an analytical paper or a lab report" (see Ambrose et al., 2010, p. 20-21).
- Sometimes learners' prior knowledge is **inaccurate**. Inaccurate prior knowledge "can be corrected fairly easily if it consists of relatively isolated ideas or beliefs that are not embedded in larger conceptual models," but misconceptions—"models or theories that are deeply embedded in students' thinking"—are extremely difficult to overcome (Ambrose et al., 2010, p. 24; see also p. 23-27). As one expert put it, "unlearning something or substituting new learning for old learning is easily twice as difficult as learning it the first time" (Richard Clark, interview, 2014).

Prior knowledge that is inactive, insufficient, inappropriate, or inaccurate hinders learning. Diagnosing precisely what a student already knows about a topic is complex – many edtech products already attempt to do this with mixed success. There are strategies (such as inserting a question that asks students to consider what they know about the topic before they begin a learning task) that can help students activate their prior knowledge and serve as the basis for discussions that surface misconceptions and other issues (Ambrose et al., 2010, p. 31), but more work remains to be done for technology to truly understand and leverage students' prior knowledge.



FIGURE 2: Not all prior knowledge is created equal. Prior knowledge must be activated, sufficient, appropriate, and accurate in order to help learning (diagram taken from Ambrose et al., 2010, p. 14).

REVIEW: Prior knowledge is stored in chunks that can be pulled seamlessly into working memory – there, it is used as a lens for examining new ideas and as Velcro that helps them stick. However, prior knowledge needs to be activated, sufficient, appropriate, and accurate in order to fully support a learning objective.

For readers who wish to know more about activating and assessing students' prior knowledge and addressing their misconceptions, we highly recommend Ambrose et al., 2010, Ch. 1.

3 Organize knowledge

PREVIEW: novices store information in sparse, superficial knowledge structures – as part of developing greater expertise, they need to begin weaving their knowledge into coherent structures based on abstract principles.

Integrating new information and prior knowledge into **coherent structures** based on **abstract principles** is essential for creating expertise. Experts don't just have more prior knowledge than novices – they organize that knowledge around meaningful principles and abstract concepts that are easy to use and apply in new contexts. This helps because “understanding new ideas is mostly a matter of getting the right *old* ideas into working memory and then rearranging them – making comparisons we hadn't made before, or thinking about a feature we had previously ignored” (Willingham, 2009, p. 91). The way experts organize their knowledge “can facilitate learning, performance, and retention” (Ambrose et al., 2010, p. 43-44).

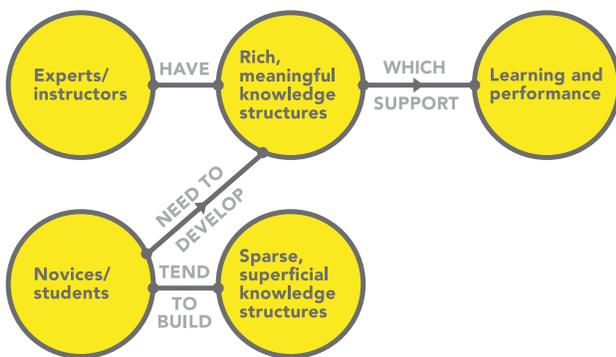


FIGURE 3: Experts organize knowledge differently than novices, and the rich knowledge structures they build improve their performance when it comes time for them to use their knowledge (diagram taken from Ambrose et al., 2010, p. 45).

All learners tend to build associations between objects that are visually similar (ex: a basketball and a globe), ideas that share meaning (ex: fairness and equality), and events that are linked with a causal relationship (ex: flipping a switch and a light turning on) (Ambrose et al., 2010, p. 46) – they do this unconsciously and effortlessly. But the connections that novices make between ideas are fewer in number and density than the connections that

experts make within a given domain. This means that all too often, important pieces of information that novices have “learned” about a subject end up being isolated and inaccessible to them later on.

Experts “have the benefit of flexibly using multiple knowledge organizations” (Ambrose, 2010, p. 56). For example, a “historian could draw on his or her knowledge in a way that is organized around theories, methodologies, time periods, topic areas, or historical figures, or combinations of these” (Ambrose, 2010, p. 57, emphasis added). This gives experts multiple ways to access any given piece of information and helps them see and instinctively respond to patterns. Experts’ deep knowledge allows them to see not just the parts, but also the whole (Willingham, 2009, p. 95).

REVIEW: Novices store information in sparse, superficial knowledge structures – as part of developing greater expertise, they need to begin weaving their knowledge into coherent structures based on abstract principles.

For readers who wish to know more about how to help students organize their knowledge more effectively, we highly recommend Ambrose et al., 2010, Ch. 2.

4 Manage cognitive load

PREVIEW: Essential and generative processing are critical for learning, but both of them require working memory. This makes it important to design instructional tasks that require very little extraneous processing.

Recall that while long-term memory can store a large amount of knowledge, working memory—the part of our mind that we use to think and actively process new information—has a severely constrained capacity in comparison. Pop science and folk wisdom has taught that most people can hold “seven plus or minus two” things in working memory at a given time (Clark and Mayer, 2011, p. 37); more recent research suggests that the real number is likely closer to four things and that

younger children have even less space in working memory than adults (Hess and Saxberg, 2013, p. 40; Richard Clark, interview, 2014).

When learners' working memory is overloaded (or underloaded) their performance degrades (Ambrose et al., 2010, p. 106). Since learners can't increase the size of their working memory (Willingham, 2009, p. 114), it's important to use other strategies to manage learners' **cognitive load**, or the amount they are asked to keep in mind at any one time.

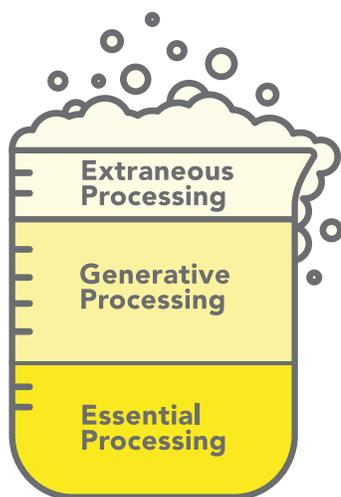


FIGURE 4: There are limits to working memory. Learners need space in working memory for essential and generative processing, which means minimizing the extraneous processing a task requires.

There are three kinds of cognitive demands that can take up space in learners' working memory, all of which limit the amount a learner can process at any one time:

- **Essential processing** is "cognitive processing aimed at mentally representing the core material (consisting mainly of selecting the relevant material)" (Clark and Mayer, 2011, p. 37).
- **Generative processing** is "cognitive processing aimed at deeper understanding of the core material (consisting mainly of organizing and

integrating)" – this is processing that allows the learner to make sense of the material (Clark and Mayer, 2011, p. 37).

- **Extraneous processing** is "cognitive processing that does not support the instructional objective" – in the context of instructional software, it "is created by poor instructional layout (such as having a lot of extraneous text and pictures)" (Clark and Mayer, 2011, p. 37).

Essential and generative processing are required for learning – learners must be able to identify new content that is important, then organize and integrate that content with their prior knowledge. However, it's possible to have too much of a good thing (to require too much essential processing) – learning tasks that introduce too much content, or content that is too complex, will overload a learner's working memory. This makes it impossible for the learner to engage in sufficient generative processing and means that new content won't be cemented in long-term memory.

Segmenting, or breaking lessons into manageable parts, is a particularly powerful method for managing the essential processing that a task requires. For example, segmenting a procedure and presenting just a few steps at a time helps learners tackle more difficult material and gives them the room they need in working memory to actively process the new information (Clark and Mayer, 2011, p. 209).

There's a large temptation for developers of instructional software to add glitz and shine as they try to compete with other kinds of media for students' attention. However, research has repeatedly shown that adding gratuitous extras (animated characters, background music, etc.) can have a negative impact on student learning (Clark and Mayer, 2011). Ruth Clark and Rich Mayer have created a set of instructional design principles that can help guide developers as they seek to decrease the extraneous processing that their tools require.

Some highlights:

- The **multimedia principle**: use words and graphics rather than words alone.
- The **contiguity principle**: align words to corresponding graphics.
- The **modality principle**: present words as audio narration rather than on-screen text.
- The **redundancy principle**: explain visuals with words in audio OR text: not both.
- The **coherence principle**: leave out material that doesn't support the instructional goal.

For more on each of these, we recommend Clark and Mayer's excellent book, *E-Learning and the Science of Instruction*.

Good instructional design can manage the amount of essential processing that a task requires, allow for sufficient generative processing, and minimize the amount of extraneous processing that a task demands.

All this to say: instructional design is complex and while there are a few cases where designers may stumble into a format that aligns with research-based design principles (take, for example, Khan Academy) the majority of designers need training on how to manage a learner's cognitive load so she can actively process the task at hand. Luckily, while there are important research questions that remain, considerable work has already been done in this area and is ready to be applied in edtech products – a key issue at the moment is finding ways to diffuse this knowledge.

REVIEW: Essential and generative processing are critical for learning, but both of them require working memory. This makes it important to design instructional tasks that require very little extraneous processing.

For readers who wish to know more about how to reduce the extraneous processing that a task requires of learners, we highly recommend Clark & Mayer, 2011, Ch. 2 (and all).

CHALLENGE	DESCRIPTION	SOLUTION	EXAMPLES
Too much essential processing	The content is so complex that it exceeds mental capacity	Use techniques to reduce content complexity	<ul style="list-style-type: none"> • Segment content into small chunks • Use pretraining to teach concepts and facts separately
Insufficient generative processing	The learner does not engage in sufficient processing to result in learning	Incorporate techniques that promote psychological engagement	<ul style="list-style-type: none"> • Add practice activities • Add relevant visuals
Too much extraneous processing	The mental load caused by extraneous and essential processing exceeds mental capacity	Incorporate techniques that promote psychological engagement	<ul style="list-style-type: none"> • Add practice activities • Add relevant visuals

TABLE 2: Adapted from Clark & Mayer, 2011, p. 38.

5 Practice deliberately

PREVIEW: Deliberate practice helps learners automate new knowledge and skills. In order to be effective, practice should incorporate goal-setting, be set at an appropriate level of challenge, come in sufficient quantity, and be spaced over time.

As we've discussed, experts have more domain knowledge than novices do, and they also store their knowledge in complex structures that are conducive to retrieving and reorganizing information when it's needed.

But another factor distinguishes experts from novices: this is the degree to which experts have **automated** key procedures, gaining the ability to take action and make decisions almost without thinking about it. This automation is a result of experts' having transferred procedural knowledge into their long-term memory, and it comes only after deliberate, repeated, effortful practice – the "only path to expertise, as far as anyone knows, is practice" (Willingham, 2009, p. 137).

“Ideological debates about learning have obscured the reality that students need enough repeated practice and feedback to build long-term memory, if they are to free room for working memory to tackle new tasks and challenges”

Hess and Saxberg, 2013, p. 43

Automatic processes “require little or no working memory capacity” – while working memory is currently thought to be “more or less fixed,” practicing skills to the point of automaticity is like chunking insofar as it serves as a cheat for increasing the amount a learner can process in her working memory (Willingham, 2009, p. 110). Practice “reinforces the basic skills that are required for the learning of more advanced skills, it protects against forgetting, and it improves transfer” (Willingham, 2009, p. 108). To these ends, considerable practice is essential for both novice and expert learners.

Deliberate practice is a specific term used in the learning science literature, and refers to practice that:

- Focuses on a specific goal or criterion for performance,
- Targets an appropriate level of challenge relative to students' current performance, and
- Is of sufficient quantity and frequency to meet the performance criteria

Ambrose et al., 2010, p. 127

Goal-setting provides focus for learners and allows them to monitor progress, adjusting strategies when needed in order to accelerate progress. It helps learners be strategic about where to direct their energy. For example, a piano player who hasn't set a specific practice goal might rehearse a piece in its entirety, not paying special attention to a difficult section – this student would be likely to spend a high proportion of his practice time on parts of the piece that didn't need work. On the other hand, a more experienced musician would be likely to identify the most troublesome measures and drill them in isolation before integrating them into the piece as a whole – by taking this approach, she would give more of her practice time to the portions of the piece that needed the most attention and would be likely to improve more quickly. In the area of reading comprehension, research has found that “students who [have] more specific goals when they [are] learning from a text [pay] more attention to passages that [are] relevant to their goals and

hence [learn] those passages better” (Ambrose et al., 2010, p. 128). With writing, students “are more likely to apply a newly learned writing strategy if they identify where it can be used, set a goal to use it there, and assess if their use of the strategy was effective” (Graham, 2008, p. 12).

HELPS Learning

[Yellow bar]						
✓		✓		✓		
✓		✓		✓	✗	

HINDERS Learning

[Yellow bar]						
				✓	✗	

FIGURE 5: To automate new knowledge and skills, learners should set goals and engage in practice that is appropriately challenging, of sufficient quantity, and spaced over time.

When discussing practice, some authors describe an **appropriate level of challenge** with a nod to Goldilocks; that is, they advocate ‘just right’ practice that’s not too easy and not too hard (Koedinger et al., 2013, p. 936). Presenting learners with a challenge is important. There is an inverse relationship “between the ease of retrieval practice and the power of that practice to entrench learning” – in other words, when students aren’t expending effort, they’re not learning (Brown et al., 2014, p. 79). This means that learning is hard work. Motivation and mindsets aren’t the focus of this document, but they’re critical to consider as part of instructional design: students need to see value in meeting the learning goals for a given task and they need to believe themselves capable of meeting those goals before they’ll expend the effort it takes for them to learn.

Giving learners ‘just right’ practice may seem straightforward in theory, but looking at the rudimentary systems that are in place for

differentiation in most large classrooms or the mediocre quality of most adaptive software engines reveals that there are substantial challenges when it comes to implementing this idea. It’s not easy to create an individualized program that gives each student practice at his or her optimal level. (For a further discussion on this topic, see Principle #7: gradually remove scaffolds.)

In order to automate skills, a learner must have **sufficient practice**: that is, she must have enough practice in the short term to master a skill, and she must rehearse that skill over the longer term in order to maintain mastery. Both instructors and students tend to underestimate the need for practice (Ambrose et al., 2010, p. 134).

One way to make practice more effective is to alter the dosage and frequency by spacing it out. Spacing may feel unnecessary or undesirable to learners, as compared with “massed” practice – studies have shown that even in cases where spaced practice led to greater learning gains, learners believed massed practice was more effective (Brown et al., 2014, p. 47). (This may partially explain why so many of us crammed for exams when we were in school!). However, the research is conclusive: “delayed subsequent retrieval practice is more potent for reinforcing retention than immediate practice, because delayed retrieval requires more effort” (Brown et al., 2014, p. 43). Learners might be surprised to learn that “you can get away with less practice if you space it out than if you bunch it together” (Willingham, 2009, p. 120).

Lastly, **continued practice** is essential for retention. One relevant study measured the performance of over 1,000 people on a basic algebra test: some participants had completed their last algebra course as little as one month before, while other participants hadn’t taken an algebra course in 55 years. In addition to varying in age, participants varied in their math backgrounds: some had only taken a single algebra course, while others had taken multiple algebra courses and/or more advanced math courses. Not surprisingly, participants who had only taken algebra tended to perform poorly on the test if many years had elapsed since they had last

completed an algebra course. But—and this next part is particularly informative for our purposes—participants who had taken multiple math courses as much as 55 years before the test, especially those who had taken advanced math courses, performed quite well on it. To put it another way:

A student who gets a C in his first algebra course but goes on to take several more math courses will remember his algebra (even 55 years later!), whereas a student who gets an A in his algebra course but doesn't take more math will forget it. That's because taking more math courses guarantees that you will continue to think about and practice basic algebra. If you practice algebra enough, you will effectively never forget it.

Willingham, 2009, p. 118

REVIEW: Deliberate practice helps learners automate new knowledge and skills. In order to be effective, practice should incorporate goal-setting, be set at an appropriate level of challenge, come in sufficient quantity, and be spaced over time.

For readers who wish to know more about deliberate practice, we highly recommend Willingham, 2009, Ch. 5; Ambrose et al., 2010, Ch. 5; and Brown et al., 2014, Chs. 2, 3, 7, and 8.

6 Get feedback

PREVIEW: learners need specific feedback that comes early and often, and they also need opportunities to apply it. Feedback about the learner's process, about the task, and about self-regulation are all helpful; feedback about the learner as a person is not.

Detailed and actionable feedback serves to fix poor practices and reinforce good practices. **Feedback** can be defined as "information provided by an agent (e.g., teacher, peer, book, parent, self, experience) regarding aspects of one's performance or understanding" (Hattie and Timperley, 2007, p. 81). Importantly, feedback must be linked with practice – the more closely it's linked, the better.

Thus we refer to a "feedback loop" in which practice provides the content and context for feedback and feedback informs the practice itself, helping the learner to improve his performance. Deliberate practice coupled with targeted feedback enhances the quality of student learning (Ambrose et al., 2010, p. 121).

This kind of feedback which "informs students' subsequent learning" is called **formative feedback** and is distinct from **summative feedback** that "gives a final judgment or evaluation of proficiency, such as grades or scores" (Ambrose et al., 2010, p. 139). Summative feedback plays an important role in most educational systems, but is not the focus of our discussion.



FIGURE 6: When practice and feedback are part of a tight loop, the learner can use feedback to guide her practice (diagram taken from Ambrose et al., 2010, p. 126).

The **timing of feedback** is important: typically, providing feedback early and often is better than the alternative. Giving feedback early prevents learners from developing misconceptions and cementing errors in their long-term memory. GPS systems provide a helpful analogy here: one of the key features of these devices is that they give corrective feedback "when the driver needs it to support the goal of reaching a particular destination as quickly as possible" (Ambrose et al., 2010, p. 142).

To be optimally useful, it is also important for feedback to be **specific** – it should identify “particular aspects of students’ performance they need to improve rather than providing a generic evaluation of performance, such as a grade or abstract praise or discouragement” (Ambrose et al., 2010, p. 139). Specific feedback provides information on whether the learner is going in the right direction, whether he is on target to meet the goal, and on where he ought to go next (Hattie and Timperley, 2007, p. 88).

Specific feedback can occur at several levels:

- Very helpful for learning: **feedback about the process** (FP) or strategy that a learner is using “appears to be more effective than feedback [about the task] for enhancing deeper learning” (Hattie and Timperley, 2007, p. 93) – this may be because it is “a direct and powerful way of shaping an individual’s task strategy” (Early et al., 1990, as quoted in Hattie and Timperley, 2007, p. 93).
- Helpful for learning: **feedback about the task** (FT) tells the learner whether she is meeting performance standards. This kind of feedback is “most powerful when it is about faulty interpretations, not lack of information. If students lack necessary knowledge, further instruction is more powerful than feedback information” (Hattie and Timperley, 2007, p. 91).
- Helpful for learning: **feedback about self-regulation** (FR) “addresses the way students monitor, direct, and regulate actions toward the learning goal” (Hattie and Timperley, 2007, p. 93). Providing feedback about student effort may be particularly helpful, especially in the early stages of learning when students need to expend the most effort in order to succeed (Hattie and Timperley, 2007, p. 95-96).
- Not helpful for learning: **feedback about the self as a person** (FS) is not as effective as the other kinds of feedback listed above, but it “is often present in class situations and is too

often used instead of FT, FP, or FR” (Hattie and Timperley, 2007, p. 96). This kind of feedback—exemplified in comments like “good boy” or “smart girl”—“usually contains little task-related information and is rarely converted into more engagement, commitment to the learning goals, enhanced self-efficacy, or understanding about the task” (Hattie and Timperley, 2007, p. 96; see also the growth mindset-related work of Carol Dweck and others).

It’s important for the person or system providing feedback to **moderate the amount** of feedback that a learner gets at any one time. A writing teacher might be tempted to cover a poorly written essay in red ink, but research shows that making “too many comments in the form of margin notes... is often counterproductive because students are either overwhelmed by the number of items to consider or because they focus their revision on a subset of the comments that involve detailed, easy-to-fix elements rather than more important conceptual or structure changes” (Ambrose et al., 2010, p. 140). In other words, it’s wise for those providing feedback to pick their battles and choose a small number of top-priority issues to highlight for the student.

Although it may seem obvious, it’s important to note that learners must have an **opportunity to apply feedback** in order to learn from it. This can be forgotten in the day-to-day classroom context: for example, a professor wishing to expose students to multiple genres of writing might assign three major papers over the course of a semester, such as a policy briefing, a persuasive memo, and an editorial. But if he only provides notes on the final essays that are submitted, students may well struggle to apply his feedback from one paper as they go to write the next one in a different genre (Ambrose et al., 2010, p. 121, p. 141).

REVIEW: learners need specific feedback that comes early and often, and they also need opportunities to apply it. Feedback about the learner’s process, about the task, and about self-regulation are all helpful; feedback about the learner as a person isn’t.

For readers who wish to know more about corrective feedback, we highly recommend Ambrose, 2012, Ch. 5 and Hattie and Timperley, 2007.

7 Gradually remove scaffolds

PREVIEW: Learners do best when working within their zone of proximal development, but it's important for learning tasks to get more challenging over time. This means building in supports and scaffolds, then slowly fading them as learners gain more knowledge and skills.

Research has shown that “adding structure and support—also called **instructional scaffolding**—to a practice activity in or out of class promotes learning when it helps students practice the target skills at an appropriate level of challenge” (Ambrose et al., 2010, p. 132). Learners benefit from instruction and practice that is appropriate for the level of knowledge and skill they have acquired. As a learner moves along the path from novice to expert, the balance of learning activities should shift from fully guided instruction towards more complex, less guided practice.

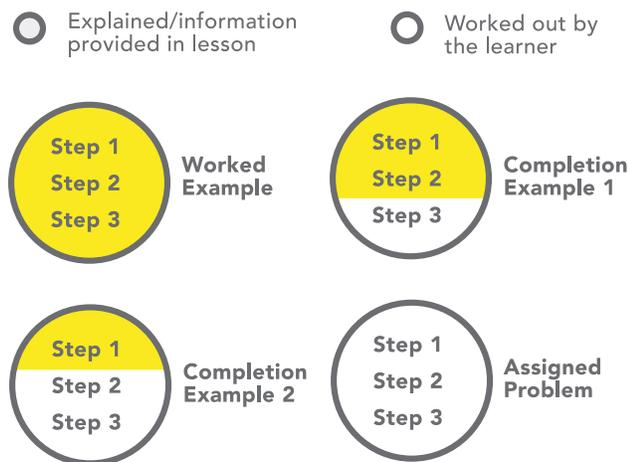


FIGURE 7: The amount explained in each example gradually fades over time (diagram taken from Clark and Mayer, 2011, p. 229).

The theoretical underpinnings for instructional scaffolding lie in Vygotsky’s notion of a **zone of proximal development**, which “defines the optimal level of challenge for a student’s learning in terms of a task that the student cannot perform

successfully on his or her own but could perform successfully with some help from another person or group” (Ambrose et al., 2010, p. 132).

In an ideal learning scenario, there is a shift over time from explicit instruction to structured practice to activities where the learner applies new knowledge in a wide array of contexts with increasing independence. The role of the instructor fades over time.

As learners develop more expertise, the goal for learning activities typically shifts from gaining knowledge to being able to articulate it and apply it in context. When learners “externalize and articulate their developing knowledge, they learn more effectively” (Bransford et al., 2000, as cited in Sawyer, 2014, p. 9). When they practice a skill repeatedly in different settings, that increases transferability (Richard Clark, interview, 2014). When they apply their knowledge to complex problems, they expand and reinforce deep structures in long-term memory that enhance task performance and improve retention. Deep learning requires that learners evaluate new ideas and relate them to conclusions. For an extended discussion of how this plays out in the context of language and literacy, see the literacy section of this document.

But learners aren’t ready for application and articulation at the beginning of a learning task. They need to start with supportive scaffolds in place, and those scaffolds need to be removed gradually over time as students are ready for ever-increasing levels of challenge.

REVIEW: Learners do best when working within their zone of proximal development, but it's important for learning tasks to get more challenging over time. This means building in supports and scaffolds, then slowly fading them as learners gain more knowledge and skills.

For readers who wish to know more about scaffolding, we recommend Reiser and Tabak, 2014.

8 Use metacognition

PREVIEW: Novices have to be taught how to use strategies to monitor and control their own learning. They need practice assessing the task at hand, evaluating their own strengths and weaknesses, planning an appropriate approach, reflecting on how things are going, and adjusting their approach when needed.

Metacognition is “the process of reflecting on and directing one’s own thinking” (National Research Council, 2001, p. 78, as quoted in Ambrose et al., 2010, p. 190). Metacognitive skills are critical for self-directed, lifelong learners, and they “arguably become more and more important at higher levels of education and in professional tasks” (Ambrose et al., 2010, p. 190). Learners who are metacognitive recognize (perhaps at an unconscious level) many of the principles that have been discussed in this paper and they are able to apply them in their own work. To take a college-level example:

Students in college are often required to complete larger, longer-term projects [than high school students] and must do so rather independently. Such projects often demand that students recognize what they already know that is relevant to completing the project, identify what they still need to learn, plan an approach to learn that material independently, potentially refine the scope of the project so they can realistically accomplish it, and monitor and adjust their approach along the way. Given all this, it is not surprising that one of the major intellectual challenges that students face upon entering college is managing their own learning.

Pascarella & Terenzini, 2005, as cited in Ambrose et al., 2010, p. 191

While various models of metacognition have been proposed by learning scientists, they all center on learners being able to monitor and control their own learning (Ambrose et al., 2010, p. 193).

To be more specific, metacognitive learners must:

- 1. Assess the task at hand.** Students need to understand what the task is asking... and they aren’t always good at listening to instructions. Even when they are listening, they may have previous experience that leads them to make incorrect assumptions about what knowledge and skills are required to complete the task, as well as what the criteria for performance will be. Learners need practice “incorporating this step into their planning before it will become a habit” and it can be helpful for them to get feedback on the “accuracy of their task assessment before they begin working on a given task” (Ambrose et al., 2010, p. 95).
- 2. Evaluate strengths and weaknesses.** Once students have a sense of what a task will require, they need to evaluate their ability to complete it. Take writing assignments as an example: students need to gauge whether they have sufficient knowledge about the assigned topic; they also need to evaluate whether they have the skills to write in the assigned genre in a way that will be compelling for their target audience. This is tricky: all people tend to have difficulty recognizing their own strengths and weaknesses, and this is especially true for students.

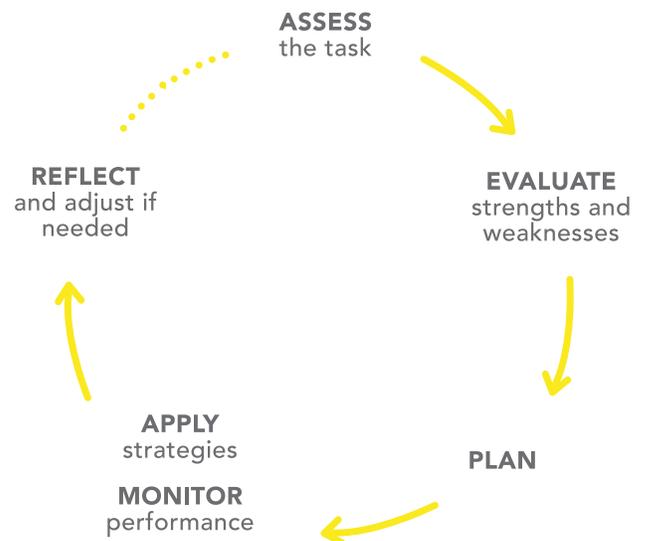


FIGURE 8: This example of a self-directed learning cycle shows the skills that an independent learner uses when she approaches a task (diagram taken from Ambrose et al., 2010, p. 193).

Moreover:

Research suggests that students with weaker knowledge and skills are less able to assess their abilities than students with stronger skills. For example, when asked to predict their performance both before and after completing a test, students in an undergraduate psychology course showed different levels of accuracy in their estimates, based on their actual performance: The highest-performing students were accurate in their predictions and postdictions (and became more accurate over subsequent tests), but the poor students grossly overestimated their performance both before and after taking the test and showed little improvement in their estimates over time.

Hacker et al., 2000, as cited in
Ambrose et al., 2010, p. 195

This goes a step further with a phenomenon called the “hard-easy effect,” which shows that learners “tend to be more accurate but under-confident on easy items, and less accurate but overconfident on difficult items” (Winne and Azevedo, 2014, p. 69). The level of knowledge learners have shapes how they approach a task, and it also shapes their motivation for carrying it out.

3. Plan an appropriate approach. Experts solve problems more quickly and accurately than novices, and they also use their work time differently: experts spend “proportionately much more time than novices planning their approach” (Chi et al., 1989, as cited in Ambrose et al., 2010, p. 195). Novices are more likely than experts to skip the planning step, taking trial and error approaches that lead to false starts. When they do make plans, those plans are less likely to be matched to the task at hand (Ambrose et al., 2010, p. 197).

4. Apply strategies and monitor performance. Learners who engage in self-explanation (talking themselves through the steps they are taking) and self-monitoring are more likely to show

learning gains, and these are skills that can be taught. For example, research has shown that “when students are taught to ask each other a series of comprehension-monitoring questions during reading, they learn to self-monitor more often and hence learn more from what they read” (Palinscar and Brown, 1984, as cited in Ambrose et al., 2010, p. 199).

5. Reflect on and adjust approach. Even when novices “monitor their performance and identify failures or shortcomings in their approach, there is no guarantee that they will adjust or try more effective alternatives” – they may be resistant to change, may lack alternative strategies, and/or may perceive switching costs to be too high (Ambrose et al., 2010, p. 199). Experts and good problem solvers, however, are likely to “try new strategies if their current strategy is not working” (National Research Council, 2001, as cited in Ambrose et al., 2010, p. 199).

Thankfully, students can be taught to become more metacognitive through a combination of explicit instruction and teacher modeling, lots of deliberate practice, and corrective feedback. “One of the most central topics in learning sciences research is how to support students in educationally beneficial reflection” (Sawyer, 2014, p. 10-11) – exploring ways to support students in building metacognitive skills could have a very large impact on the design of many edtech products.

REVIEW: Novices have to be taught how to use strategies to monitor and control their own learning. They need practice assessing the task at hand, evaluating their own strengths and weaknesses, planning an appropriate approach, reflecting on how things are going, and adjusting their approach when needed.

For readers who wish to know more about how to strengthen learners’ metacognitive skills, we highly recommend Ambrose et al., 2010, Ch. 7.

Learning to be literate

► What are the language and literacy skills that students develop in grades 4-13+?

We'll use the Common Core State Standards (CCSS) to guide our discussion of K-12 content. The standards are separated into bands for K-5, 6-8, 9-10, and 11-12. Broadly speaking, the standards cover reading, writing, speaking, and listening. Beginning in grades 9-10, there are also specific standards for literacy in history/social studies, as well as literacy in science and technical subjects.

One major shift is the increased emphasis on building knowledge through reading content-rich nonfiction. Students are expected to have regular practice reading complex texts and using the academic language that's embedded within them. They need to grow their vocabulary and content knowledge through a mix of reading, conversation, and direct instruction; they also need to use context to determine word meanings and appreciate nuances in meaning.

A second important shift is that with the CCSS, students are expected to ground their thinking in evidence from the text: this means drawing text-based inferences and citing appropriate evidence in both conversations and written work. Students

must use evidence to inform others and build persuasive arguments.

Literacy standards for higher education are significantly more fragmented than they are for K-12. While the Common Core has set goals and expectations in reading, writing, listening, and speaking, the expectations are somewhat less clear when looking across colleges and universities. Very often, "the content of college courses...is developed with no awareness of K-12 expectations or even those of other college-level courses" (Tepe, 2014, p. 22).

However, developmental English programs present an area of opportunity at the university level, as they are beginning to move toward more formal objectives and an integrated approach for learning English. For example, the Virginia Community College System has identified the following desired student learning outcomes in their developmental English curriculum guide:

- Demonstrate the use of pre-reading, reading, and post-reading skills with college-level texts.
- Pre-write, draft, revise, edit, and proofread college-level texts.

- Expand vocabulary by using various methods.
- Demonstrate comprehension by identifying rhetorical strategies and applying them to college-level texts.
- Analyze college-level texts for stated or implied main idea and major and minor supporting details.
- Demonstrate critical thinking skills when reading and writing college-level texts.
- Write well-developed, coherent, and unified college-level texts, including paragraphs and essays.
- Identify, evaluate, integrate, and document sources properly.

NROC, 2013, p. 41

These kinds of course goals overlap with the Common Core expectations for literacy and will serve to guide our discussion on post-secondary content.

The next few pages will examine the cognitive behaviors that students need to employ in order to meet expectations for language and literacy development – in other words, we’ll spend some time examining what good readers and communicators do.

► What do good readers do?

Good readers set goals and monitor progress

Setting a purpose for reading helps the reader: the reader’s goals determine what level of comprehension is going to be required, as well as the degree to which specific details need to be recalled, evidence needs to be cited, etc. Take an example: a good reader would likely approach *To Kill a Mockingbird* casually when reading it for pleasure; she’d be more likely to annotate and explicitly identify the themes of the book if she knew she had to write a critical essay based on the text. Another reader might know that scanning

a scientific article to brush up on the topic was sufficient before having dinner with a colleague; he’d read the same article with a different lens if he knew he were going to be asked to compare his research findings with those of the author.

Defining a goal “can serve as a benchmark for judging whether the reader’s standards for comprehension are satisfied” (McNamara, 2012, p. 472). Once reading goals are established, the reader must monitor whether she’s comprehending well enough to use the information from the text “to accomplish foreseeable goals (e.g., essay writing, answering questions, participating in a discussion, etc.)” (McNamara, 2012, p. 469).

Good readers prepare for reading

In addition to setting goals, good readers often preview sections of a text before they begin to read. For example, before starting a new novel, a reader might flip through the book to see how long the first chapter was so that he knew whether he was likely to reach a good stopping point before dinner. Another reader might open her biology textbook and examine the diagrams and a few terms from the glossary before beginning to read the text itself; she might also make a judgment about how much of the content seemed new and potentially difficult to understand.

This helps comprehension: previewing the text before reading can help the reader conceptualize “what is already known and, consequently, what requires further clarification when the text is read in detail” (McNamara, 2012, p. 472).

Good readers activate prior knowledge

Activating prior knowledge (Principle #2) is important to mention because struggling readers “often pay little attention to whether the text makes sense at the global level” (i.e. whether individual sentences are adding up to something larger) or whether the “overall text content is consistent with their pre-existing beliefs or understanding about the subject matter” (McNamara, 2012, p. 469-470).

In cases where a reader isn't consistently using prior knowledge as a check to see whether his understanding of the text makes sense, he's likely to "develop an illusion of understanding... one that does not accurately reflect the potentially fragmented nature of [his] mental representation" (McNamara, 2012, p. 470).

Good readers integrate concepts

This is highly aligned with Principle #3 (organize knowledge). Suffice it to say that good readers make links among ideas, synthesize ideas, and integrate concepts into a global understanding (McNamara, 2012, p. 469).

Self-explanation is a strategy that "tends to engage readers into a mode of active processing... [and helps them] detect problems in comprehension by externalizing internal thought processes... Self-explanation also helps readers to engage in elaboration, which involves relating the text content to what one already knows" (McNamara, 2012, p. 484).

Good readers self-regulate and use strategies to help when comprehension breaks down

Self-regulation is part of being metacognitive (Principle #9). Good readers notice when what they are reading stops making sense. This can happen at the local level or a more global level – a reader might realize that an unusual vocabulary word is critical for his understanding of the text and stop to look it up, or she might have to go back and re-read an entire section of *Ulysses* several times before she grasps what Joyce is trying to say. Many readers experience difficulty in constructing a "textbase understanding," or an understanding of what's explicitly stated in the text, "particularly when they encounter unfamiliar words and/or complex syntactic structure" (McNamara, 2012, p. 477).

The important thing is that "the reader should not leave comprehension problems unresolved; instead, he or she should use available sources to solve problems to attain the level of comprehension required for the circumstances"

(McNamara, 2012, p. 486).

Good readers have a toolkit of available strategies for getting unstuck – they're likely to have several frequent go-to strategies, and they're also likely to have a few more specific approaches they can try if their first attempts fail. Teaching this hasn't been straightforward. One expert described the fragmented way reading comprehension skills are commonly taught as the "skill du jour" method, or as teaching "splinter skills." She said:

When you start teaching all these splinter skills initially, it's just overwhelming. It's overwhelming to the very kids who most need to do a few things well... If we work too hard to teach all these splinter strategies, we get students who are really confused about what to do when they see print... It's like if I gave you a tool box with fifty tools in it and I said, 'Figure out which one of these will help you take the faucet off the sink.' Well, you'd be much more successful if I said to you, 'Here's three tools, they are really good tools and they fix most things for you.'

Sharon Vaughn, interview, 2015

Building proficiency with a smaller set of effective strategies might help students who have stumbled regain reading momentum while avoiding cognitive overload (Principle #4).

Good readers discuss and/or write about their reading

Good readers deepen their comprehension of the text by actively processing what they've read. Deep learning requires that learners evaluate new ideas and relate them to conclusions: "when learners externalize and articulate their developing knowledge, they learn more effectively" (Bransford et al., 2000, as cited in Sawyer, 2014, p. 9). As Graham and Hebert note in their meta-analysis *Writing to Read*, writing notes about a text, summarizing it, and responding to it can all help. (We will elaborate on discussion in the final section of this document, which is centered on argumentation).

Taking notes about a text requires learners to process information in working memory as they organize it (Principle #3) on paper. Notes can take varied forms (concept maps, structured notes, etc.) – the important takeaway is that note-taking enhances comprehension (Graham and Hebert, 2010, p. 16).

Summarizing text is one strategy that readers can use to help them organize their knowledge: it requires active cognitive processing of the text (of its content, of the argument the author is making, etc.) and deepens reading comprehension. Good readers summarize mentally as they read and they may also write summaries (Graham and Hebert, 2010, p. 15).

When students respond to a text, they might share a personal reaction, relate to their own experience, or craft an analytical interpretation of the text. Students are more likely to develop a deep understanding of the text when they “write about text in extended ways involving analysis, interpretation, or personalization,” and research has found that these activities have a strong impact on reading comprehension (Graham and Hebert, 2010, p. 13).

► **What do good writers do?**

Some of the most effective techniques for literacy instruction involve integrated approaches where reading and writing skills are developed together and used to support each other. However a good reader does not necessarily make a good writer – most teachers can think of an example of a student who can read wonderfully and can’t put together a coherent paragraph. Writing requires additional skills and strategies that can be illustrated in the three phases of writing: planning, drafting, and reviewing.

Good writers plan ahead and pre-write

Planning ahead before jumping into the writing process is essential for effective writing. The literature on writing outlines specific phases of pre-writing, such as:

- Goal setting (i.e., establishing criteria concerning how to communicate with the audience)
- Generating (i.e., retrieving relevant information from long-term memory)
- Organizing (i.e., selecting the most important information and structuring it into a writing plan)

Smagorinsky, Cambridge, 2014, p. 613)

Goal setting is a particularly important aspect of the planning process (Graham and Perin, 2007, p. 17). When students take time to understand their intended audience for a written piece and consider the tone/kind of argument that is most likely to be effective, they’re more likely to write in a style that will meet their objectives. Metacognitive writers monitor themselves throughout the writing process, referring back to their goals and making adjustments as needed (Principle #8).

Activities designed to help students activate and organize their prior knowledge (Principle #2) can be helpful in beginning the writing process and lead to better learning outcomes (Graham and Perin, 2007 p. 18). For example, instructors can encourage students to begin with a “think sheet or graphic organizer for planning a paper,” which might “direct [them] to identify their audience and purpose for writing the paper, generate possible content (in note form), decide which basic ideas to use (by putting a star next to them), and order the ideas for writing (using numbers to note what will come first, second, third, and so forth)” (Graham and Perin, 2008, p. 9).

Pre-writing can be especially helpful for the kind of long-form writing tasks that are often daunting to students. Pre-writing activities can break the difficult process of planning into component parts, helping students activate their prior knowledge and organize thoughts on paper, which opens capacity in working memory for the act of writing itself.

Good writers draft consciously and strategically

Planning and pre-writing can reduce the cognitive load placed on working memory during the writing

process (Principle #4), but can't eliminate it (Steve Graham, interview, 2015). Researchers have found that good writers tend to engage in "abundant local planning" (e.g., pausing after they finish a clause or sentence to plan the next one) while they are actively producing a text (Smagorinsky and Mayer, 2014, p. 613). In fact, most cognitive activity for writers of all ability levels "actually takes place at the sentence construction level" (Steve Graham, interview, 2015).

This is a key difference between writing and reading: while significant portions of the reading process (ex: decoding) can be automated, the process of writing cannot become fully automated and will always require conscious attention and decision-making on the part of the writer.

Metacognition is a key aspect of this phase of writing. Good writers consider who the audience is, how they're meeting their goals for a particular piece of writing, and whether their writing is going to be clear to the reader. They keep their initial goals in mind and are flexible about updating them as the writing process develops.

Good writers edit and review

Iterating is an essential part of writing. Good writers assess their own work based on the goals they set for a piece during the planning stages of writing; they also go through multiple drafts of major writing projects, adjusting substance and improving grammar, usage, and mechanics.

However, inexperienced writers are typically unable to detect most of the errors in their writing and they're unable to correct more than half the errors they do find (Smagorinsky and Mayer, 2014, p. 614). There are instructional strategies that help here – for example, teaching students to "apply reduction principles (e.g. delete unnecessary material, delete redundant material, select a word to replace similar ideas or items, draft a topic sentence)" can help them learn to approach their writing more strategically (Graham, 2008, p. 10).

But writers also benefit from getting feedback from others as part of the revision and editing process (Principle #6). Instructors often focus feedback on these features of a text:

- Are ideas in text clearly presented and fully developed?
- Is the text easy to follow and logically organized?
- Are words used effectively and precisely?
- Are sentences varied to promote fluency, rhythm, and natural speech patterns?
- Does the text capture appropriate tone or mood to make maximum impact on the reader?
- Are there spelling, usage, and grammar errors?
- Is the written product legible, attractive, and accessible?

Graham, 2008, p. 11

To the extent that writers can internalize these questions and learn to use them to assess and improve their own writing, each draft they produce can be substantially improved.

Peer feedback can be another effective form of review – it benefits both the author and the reviewer:

What we know pretty well is that if peers give each other feedback in a structured way, and they're taught how to do it and it's positive, it has a positive effect on both peers giving the feedback and peers receiving the feedback.

Steve Graham, interview, 2015

Becoming a good writer is a powerful thing – writers who communicate clearly and artfully are able to inform, persuade, and entertain their peers

and colleagues. Writers write for their readers, which places writing in a profoundly social context. And when leveraged appropriately, this can be a huge motivating factor – good news indeed for the instructional designers and teachers who need their young writers to persist throughout the years of practice that it will take to develop expertise.

► **What do good communicators do?**

Good communicators listen attentively and respond persuasively. They engage in constructive argumentation. While a narrow definition of literacy emphasizes the comprehension and production of text, students in grades 4-13+ also need to participate in collaborative discussions and use their oral language skills to learn from, inform, and persuade others. They use argumentation skills in social contexts in real time.

An emphasis on argumentation

Argumentation is one application of speaking and listening that receives substantial attention as part of the Common Core. Component skills include the abilities to:

- Propel conversations by posing and responding to questions that probe reasoning and evidence; ensure a hearing for a full range of positions on a topic or issue; clarify, verify, or challenge ideas and conclusions; and promote divergent and creative perspectives (CCSS: grades 11-12, SLS.CC.1.c).
- Evaluate a speaker’s point of view, reasoning, and use of evidence and rhetoric, assessing the stance, premises, links among ideas, word choice, points of emphasis, and tone used (SLS CC.11-12.3).

Argumentation incorporates both listening and speaking into a complex cognitive and social exercise that is essential for advancing literacy across the board:

The underlying structure of effective academic writing— and of responsible public discourse— resides not just in stating our own ideas but in listening closely to others around us, summarizing their views in a way that they will recognize, and responding with our own ideas in kind.

Graff and Birkenstein, 2006, p. 3

So all students have to do to become academics is argue more? Not exactly. Most students argue regularly – in social interactions, peers contradict each other frequently (and siblings can find it within themselves to bicker about just about anything!). Yet in more academic settings, those same students struggle to summarize one another’s perspectives and respond to one another clearly and succinctly as part of an academic discourse.

Using explicit instruction to break down the structural elements of good arguments can be illuminating for students. There are many ways to structure an argument. For example, expository arguments often use a “claim, data, warrant, backing, qualifier, and rebuttal” structure (Andriessen and Baker, 2014, p. 440-441). Dissecting the structure of an effective argument helps students organize their thoughts (Principle #3) and call on processes (argument schemas) that improve their logical argumentation performance. Getting explicit instruction in these “ground rules” improves students’ learning and collaboration skills (Andriessen and Baker, 2014, p. 442).

Using evidence is a particularly important element of argumentation, and it’s one that’s challenging for many students. The idea that “evidence is relevant to argument and essential in supporting and refuting claims” may seem obvious to adults (Kuhn, 2010, p. 814). However, students need focused practice (Principle #5) in both identifying the evidence that best supports their argument, and also in using that evidence effectively.

It’s worth noting that developing argumentation skills is related to developing social skills,

particularly when discussions are practiced in a collaborative manner with the goal of converging on a solution (Andriessen and Baker, 2014, p. 442). Through constructive argumentation, students learn to consider both sides of an issue, empathizing and engaging with others. This process of collaborative learning can help students develop the skills necessary for the group work

they do in school, as well as the teamwork they will need to do in professional settings in the future (Miyake and Kirschner, 2014, p. 421). As with writing, the social context for argumentation can be a powerful motivating factor for students – this has implications for both instructional designers and the teachers who are charged with developing students' literacy skills.



In conclusion

While learning scientists would be the first to admit that we all still have a lot more to learn about learning, there's remarkable consensus about principles that have been discussed in this document – the research that's been done in this discipline helps us understand the mental work that students need to do in order to build expertise in every content area, including literacy.

We hope this document helped accelerate your understanding and will serve as a useful reference for you and your colleagues. We also hope that we've piqued your curiosity and equipped you to go learn more. More of us who work as learning designers—including those who are on the ground with students, engaged in the difficult day-to-day work of teaching and learning—become aware of this research every day. Learning science is being used to inform lesson plans, instructional programs, school models, technology solutions, and more.

As we work together to apply these principles in a variety of settings, there's great potential for improving student outcomes.

It's tremendously exciting to imagine the ways our collective impact will grow as we begin creating solutions firmly rooted in an understanding of how students learn, and as we continue to share these lessons with those around us.

A word of thanks: Those of us here at Intentional Futures are deeply indebted to others who have helped us on our learning journey. In addition to reading widely, our team benefited from conversations with dozens of experts in the field. We're particularly grateful for the time and assistance of Richard Clark, Steve Graham, Elijah Mayfield, Eileen Murphy, Bror Saxberg, and Sharon Vaughn. We're also thankful for the partnership of Sara Allan, Jill Nishi, Jason Palmer, Tom Stritikus, and all our colleagues at the Bill & Melinda Gates Foundation.

WE INVITE YOU TO SHARE THOUGHTS ON YOUR OWN LEARNING JOURNEY

WWW.INTENTIONALFUTURES.COM TWITTER: @INTENTFUTURES

References

Alliance for Excellent Education. *Saving Now and Saving Later: How High School Reform Can Reduce the Nation's Wasted Remediation Dollars*. May, 2011. Accessed Feb. 2, 2015: <http://all4ed.org/wp-content/uploads/2013/06/SavingNowSavingLaterRemediation.pdf>

Ambrose, Susan A., Michael W. Bridges, Michele DiPietro, Marsha C. Lovett, and Marie K. Norman. *How Learning Works: Seven Research-Based Principles for Smart Teaching*. San Francisco, CA: Jossey-Bass, 2010. Print.

Andriessen, Jerry and Michael Baker. "Arguing to Learn." *The Cambridge Handbook of the Learning Sciences, 2nd Edition*. New York, NY: Cambridge UP, 2014. Print.

Annie E. Casey Foundation. *Early Warning! Why Reading by the End of Third Grade Matters*. Jan., 2010. Accessed Feb. 2, 2015: <http://www.aecf.org/resources/early-warning-why-reading-by-the-end-of-third-grade-matters/>

Annie E. Casey Foundation. *Early Reading Proficiency in the United States*. Jan., 2014. Accessed Feb. 2, 2015: <http://www.aecf.org/m/resourcedoc/aecf-EarlyReadingProficiency-2014.pdf>

Bill & Melinda Gates Foundation. *Literacy Courseware RFP*. Feb. 2013. Accessed Jan. 30, 2015: <https://docs.gatesfoundation.org/documents/Literacy-Courseware-Challenge-RFP.pdf>

Bill & Melinda Gates Foundation. *Next-Generation Courseware Challenge to Improve U.S. Undergraduate Education*. May, 2014. PDF, not publicly available.

Brown, Peter C., Henry L. Roediger, III, and Mark A. McDaniel. *Make It Stick: the Science of Successful Learning*. Cambridge, MA: Belknap Press, 2014. Print.

Clark, Richard E. *Cognitive Task Analysis for Expert-Based Instruction in Healthcare*. In Specter, J.M. Merrill, M.D. Elen, J. and Bishop, M.J. (Eds.). *Handbook of Research on Educational Communications and Technology, 4th Edition*. New York, New York: Springer. p.541-551, 2014. PDF.

Clark, Richard E., Paul A. Kirschner, and John Sweller. *Putting Students on the Path to Learning: The Case for Fully Guided Instruction*. *American Educator*, Spring 2012.

Clark, Ruth Colvin, and Richard E. Mayer. *E-Learning and the Science of Instruction: Proven Guidelines for Consumers and Designers of Multimedia Learning*. San Francisco, CA: Pfeiffer, 2011. Print.

Farrington, C.A., Roderick, M., Allensworth, E., Nagaoka, J., Keyes, T.S., Johnson, D.W., & Beechum, N.O. *Teaching adolescents to become learners: The role of noncognitive factors in shaping school performance: A critical literature review*. Chicago: University of Chicago Consortium on Chicago School Research, 2012: <https://ccsr.uchicago.edu/sites/default/files/publications/Noncognitive%20Report.pdf>

Gianneschi, Matt and Mary Fulton. *A Cure for Remedial Reporting Chaos: Why the U.S. Needs a Standard Method for Measuring Preparedness for the First Year of College*. Educational Commission of the States. June, 2014. PDF.

Graff, Gerald, and Cathy Birkenstein. *They Say, I Say: The Moves That Matter in Academic Writing*. New York, NY: Norton, 2006. PDF.

Graham, Steve. *Effective Writing Instruction for All Students*. Wisconsin Rapids, WI: Renaissance Learning, 2008. PDF.

Graham, Steve, and Michael Hebert. *Writing to Read: Evidence for How Writing Can Improve Reading – A Report from Carnegie Corporation of New York*. Washington, DC: Alliance for Excellent Education, 2010. PDF.

- Graham, Steve, and Dolores Perin. *WritingNext: Effective Strategies to Improve Writing of Adolescents in Middle and High Schools – A Report to Carnegie Corporation of New York*. Washington, DC: Alliance for Excellent Education, 2007. PDF.
- Hattie, John, and Helen Timperley. *The Power of Feedback*. *Review of Educational Research*. March 2007, Vol. 77, No. 1, p. 81-112. PDF.
- Hess, Frederick M., and Bror Saxberg. *Breakthrough Leadership in the Digital Age: Using Learning Science to Reboot Schooling*. Thousand Oaks, CA: Corwin, 2014. Print.
- Kalyuga, Slava, Paul Chandler, and John Sweller. *Managing Split-attention and Redundancy in Multimedia Instruction*. *Applied Cognitive Psychology*. Vol. 13, p. 351-371, 1999. PDF.
- Koedinger, Kenneth R., Julie L. Booth, and David Klahr. *Instructional Complexity and the Science to Constrain it*. *Science*, Vol. 342, Nov 2013. PDF.
- Kuhn, Deanna. *Teaching and Learning Science as Argument*. *Science Education*, 94, p. 814-824, 2010. PDF.
- McNamara, Danielle S. *Reading Comprehension Strategies: Theories, Interventions, and Technologies*. Psychology Press, 2012. Print.
- Miyake, Naomi and Paul A. Kirschner. "The Social and Interactive Dimensions of Collaborative Learning." *The Cambridge Handbook of the Learning Sciences, 2nd Edition*. New York, NY: Cambridge UP, 2014. Print.
- Nathan, Mitchell J. and R. Keith Sawyer. "Foundations of the Learning Sciences." *The Cambridge Handbook of the Learning Sciences, 2nd Edition*. New York, NY: Cambridge UP, 2014. Print.
- National Center for Education Statistics. *Fast Facts*. Accessed Jan. 24, 2015: <http://nces.ed.gov/fastfacts/display.asp?id=96>
- The Nation's Report Card. *Writing 2011, National Assessment of Educational Progress at Grades 8 and 12*. 2011. Accessed Jan. 30, 2015: <http://nces.ed.gov/nationsreportcard/pdf/main2011/2012470.pdf>
- The Nation's Report Card. *Are the nation's twelfth-graders making progress in mathematics and reading?* 2013. Accessed Feb. 2, 2015: http://www.nationsreportcard.gov/reading_math_g12_2013/#/
- NROC. *Developmental English – An Integrated Program; Product Design Report Appendices*. Marina, CA: Monterey Institute for Technology and Education, 2014. PDF.
- Reiser, Brian and Iris Tabak. "Scaffolding." *The Cambridge Handbook of the Learning Sciences, 2nd Edition*. New York, NY: Cambridge UP, 2014. Print.
- Sawyer, R. Keith. "Introduction: The New Science of Learning." *The Cambridge Handbook of the Learning Sciences, 2nd Edition*. New York, NY: Cambridge UP, 2014. Print.
- Smagorinsky, Peter and Richard E. Mayer. "Learning to Be Literate." *The Cambridge Handbook of the Learning Sciences, 2nd Edition*. New York, NY: Cambridge UP, 2014. Print.
- Tepe, Lindsey. *Common Core Goes to College: Building Better Connections Between High School and Higher Education*. New America, 2014. PDF.
- Willingham, Daniel T. *Why Don't Students Like School?: A Cognitive Scientist Answers Questions About How the Mind Works and What It Means for the Classroom*. San Francisco, CA: Jossey-Bass, 2009. Print.
- Winne, Philip H. and Roger Azevedo. "Metacognition." *The Cambridge Handbook of the Learning Sciences, 2nd Edition*. New York, NY: Cambridge UP, 2014. Print.