PART I

Building the foundation

Evidence-informed learning experiences

Introduction

In the first part of this book, we try to build a foundation that will get us grounded into how to design learning experiences based on research from the learning sciences. As authors, it's crystal clear to us that laying the foundation first – both in this book and in general – is the only way to successfully improve our practice of designing learning experiences. No matter how we slice it, what it comes down to is: *We need to use the evidence available to us to make sure we move beyond opinions and intuition.* Having done this, we can then effectively spend our and our organization's money, time and effort based on informed decisions, and create the best learning experiences we can for our learners and the organizations they work in. When we talk about designing learning experiences based on the learning sciences, we call this **evidence-informed practice**. We explain in the first chapter what we mean by this and why we think it's important to distinguish between an **evidence-based** and an evidence-informed approach. We also give some ideas on how to get started when designing learning experiences in an evidence-informed way. We then dive into the learning sciences and discuss why they matter to us as practitioners. After all, the learning sciences are where we can find the evidence we need to improve the learning experiences we design.

To be truthful, just as we discussed what description we would use for our roles throughout this book (learning professionals), we have also debated which term we would use in this book to describe our practice. Plenty of options, after all. We could use **instructional design**, learning design, and the more recent term **learning experience design**. We have opted for the last, although we find 'learning experience design' to be an awkward term and we really wouldn't have minded using 'instructional design' or 'learning design' instead. But we're realists and we see, in some organizations, a devaluation of the instructional or learning design role. For example, someone with the title instructional designer can in fact be a content developer. It's fairly unclear within organizations, even among us practitioners, what instructional design, learning design or even learning experience design is. We opted to use learning experience design, but we're definitely not married to the term, as long as we focus on building on a strong foundation and then designing learning experiences based on evidence from the learning sciences. But we use the term throughout this book and discuss in Chapter 2 what we mean by it. We also explain what we believe it means to design learning experiences because, no, you don't necessarily always have to *design* all learning experiences. After all, people have learning experiences all the time, be it in an informal or serendipitous way, during or outside work. Last but not least, in this first part, we reveal what it means to design three-star learning experiences.

1

Designing learning experiences in an evidence-informed way

We draw a distinction between evidence-based and evidence-informed, as there's a difference.

Evidence-based practice is an interdisciplinary approach to clinical practice, grounded in medicine. Sackett *et al* (1996) see it as a three-legged stool integrating three basic principles (Figure 1.1):

1. the best available research evidence bearing on whether and why a treatment works
2. clinical expertise of the health care professional (clinical judgement and experience) to rapidly identify each patient's unique health state and diagnosis, their individual risks and benefits of potential interventions
3. client preferences and values.

FIGURE 1.1 Evidence-based practice

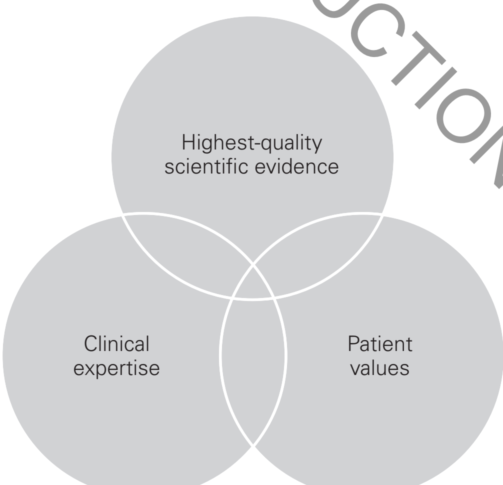

For example, if a decision is made on the dosage, intake and working of a medicine, then it means that it was tested and approved for a specific sickness or condition and a specific part of the population (eg a person aged between 30 and 50, with a healthy BMI and symptoms X, Y and Z). The instructions to take that medicine in the morning on an empty stomach allows for a wide range of specific circumstances (at home, in the car, on the beach, when and wherever as long as it's on an empty stomach in the morning). The problem is that for learning, this doesn't work.

Evidence-informed still means 'based on scientific research'. But in the field of the learning sciences we're dealing with muddy real-life things, also known as 'variables', that can influence what we want to achieve and whether it's achieved. In our context of designing a learning experience, this is an intervention's effects. Our field simply doesn't always allow for 'straightforward measurement'. We simply can't usually deliver the same quality of evidence as clinical practice does. This is because in learning environments, we're dealing with many different variables that interact and are hard to control. Literally, what worked in a workshop or lesson today will not necessarily work in the same workshop or lesson that afternoon, one day later or three months later. Just the fact that the learners are different, with different **prior knowledge**, different beliefs, different needs and/or different motivations to participate, can change everything. And then there are also environmental factors. Take the room where the instruction or training takes place, for example. If the first room has the 'right' temperature and the second is too hot or too cold, it can impact the learning experience. Or, if the first cohort is local (eg either from the same company/department and/or in the same room) and the second global (eg from different companies and/or spread throughout the globe), this will impact how the learners interact with each other and with an instructor, and, thus, how they'll learn. Hence, when we *use* evidence, we need to acknowledge that what works in one context might not necessarily work in another.

Also, we often use more **qualitative data** and this type of data provides weaker evidence than **quantitative data**. Quantitative data are numerical. They're about 'hard' numbers, measurable variables that can be used for mathematical calculations and statistical analyses. Examples are people's weight or height, or the number of times they took a medicine, or the score on an achievement test. It's important to note here that numerical doesn't mean that it's per se exact, reliable or valid. A self-report on a Likert scale (eg how much did you think you learned: 1 = very little through 5 = a lot) yields numerical data but is, of course, subjective and thus not very reliable.

Qualitative data are information about characteristics, properties or 'qualities', such as quality of the facilitator, engagement, or perceived usefulness of a training. This type of data describes but doesn't define. It approximates but doesn't measure the actual attributes (eg how much you think you learned as opposed how much you actually learned). A third problem – and this can be either quantitative or qualitative – is that we sometimes can't measure what we need to measure and then we use a **proxy**. For example, in a learning context, we often ask learners if they were engaged or we observe them and rate their engagement. When they indeed felt engaged or were observed as being engaged, we take that as a positive proxy for learning. But in reality, we haven't measured learning, we've measured either how busy they were (but with what?) and/or their perception of their own engagement. In other words, in our field, the evidence will alert us to what might work and under what conditions (ie it will inform us as to whether it might work, when it might work, how it might work etc), but this is no guarantee that it will work!

This nuanced description of the difficulties with using scientific evidence in our field might suggest that it's too difficult or not worth it. This is definitely *not* the case. Although we need to acknowledge that in our field the evidence is usually weaker than in a field like medicine (though placebos have an effect in medicine and we measure pain via subjective rating scales) or physics, this doesn't mean we shouldn't use the evidence. On the contrary. This whole book is about why we should use the evidence that's available to us. We should use it based on our practical wisdom and based on the context we work in. This combination will help us explain why we make certain decisions and it will make us better practitioners overall.

There are also similarities to medicine. Here too we can speak of a three-legged stool. But we can also use other types of evidence to decide what works best to achieve a certain goal through a learning experience. Examples are input from learners and stakeholders, data from systems that might be used in the workplace, and of course our own expertise as learning professionals (see Figure 1.2).

In this book, we focus on the scientific evidence, as this is underused when it comes to designing learning experiences and, in particular, evidence from the learning sciences: an interdisciplinary field focused on developing a deeper scientific understanding of learning. We explore its history and meaning, as well how it's useful to us as learning professionals, in the next section.

FIGURE 1.2 Evidence-informed practice for learning experience design

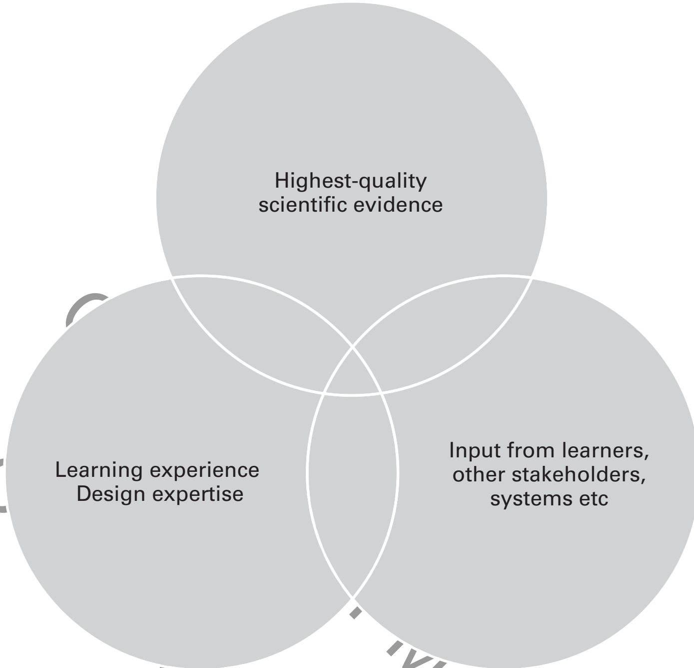

Although the tide is turning in that more professionals in our field are beginning to recognize that scientific evidence needs to be leveraged to make informed design decisions, we still too often base our designs on hunches and beliefs (and sometimes worse – based on myths!), aesthetics and learner or stakeholder opinions and preferences. This book is a strong plea for consistently integrating evidence from the learning sciences into our practice so that we can truly support the organization we work for to support their people to continuously learn so that they can do their jobs better.

We're both members of th*e* Debunker Club, founded by Will Thalheimer and dedicated to eradicating learning myths and sharing proven evidence-informed insights. We're committed to this ourselves because we agree with the club that, when we design learning experiences based on myths, we're spending time and money that could be better spent elsewhere. After all, we're there to help organizations and the people who work there. Possibly even more important from an ethical point of view, we also hurt learners when we incorporate learning myths and **misconceptions** into our designs. And last, but not least, it's also detrimental to the foundation of the learning profession when we base our designs on hype, myths, anecdotes from gurus, silliness or sexy bells and whistles. We simply must do all we can to debunk myths and misconceptions, learn from research and practice, and share evidence-informed information. We need to be open yet sceptical.

We imagine you looking at us with slight uncertainty, wondering if this means you need to dive into the learning sciences research yourself. We can reassure you that you don't. Not necessarily, anyway. There are other ways for all of us to put our research hat on to prevent ourselves from getting fooled or to confirm that something we're reading, seeing or hearing might actually be true. Daniel T Willingham's steps come to the rescue (Figure 1.3). Let's see how we can use them to our benefit.

FIGURE 1.3 Steps towards designing learning experiences in an evidence-informed way

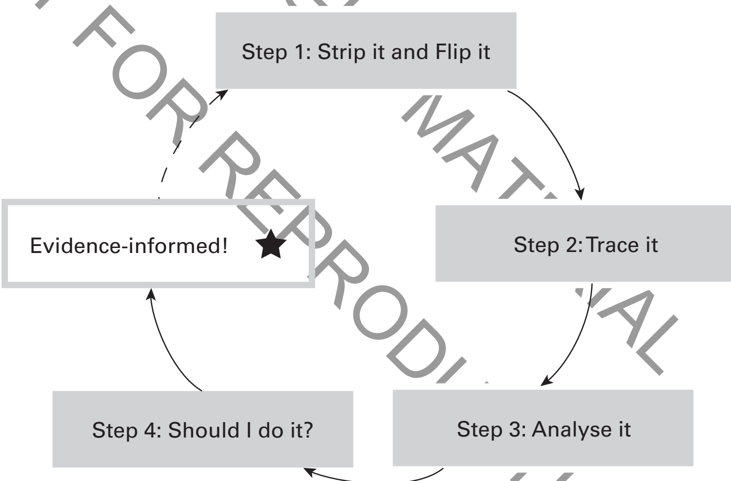

Adapted from Willingham (2012a, 2012b)

Steps to start designing learning experiences in an evidence-informed way

Let's look at an example excerpt from an article and then figure out what we need to do if we follow Willingham's steps. The following example is based on an existing article; we have adapted it for the purpose of practice and because it's about informing and not naming and shaming.

EXAMPLE How to design learning experiences for millennials
Joe Doe

Millennials are still the largest generation in the global workforce (Research Centre Y). Organizations need to ask themselves what this means for training their staff. How should the approach differ from previous approaches taken with Baby Boomers or even Gen Xers? Millennials are the future, so it's critical that organizations design learning strategies that ensure they get the most out of this unique generation.

How millennials are different

Millennials are also known as the 'digital native' generation. Because of their lifelong relationship with technology, they learn differently. Don't bother millennials with lecture-style training; it won't work. You need to focus on hands-on, experiential learning, as only then will you get the results that your organization need you to deliver.

A study by We Do Research, comparing the learning preferences of 5,000 people in one organization, clearly shows that millennials want to spend 10 per cent less time on in-person training. Evidence clearly shows that millennials prefer authentic learning contexts, enabling them to connect learning to their jobs.

What works best for millennials?

There's only one simple answer to this question: video. It is the only medium that grabs a millennial's attention. The average human brain processes video 60,000 times faster than text and a millennial brain even processes it 100,000 times faster.

A study by ABC clearly shows, in addition to the fact that video is way more enjoyable than reading a book or listening to someone talking, that learning through video is 80 per cent more effective for millennials than any other medium.

And so, it continues....

Step 1a: Strip it

The first part, 'strip it', means that we take a critical look at the language used. In the case of the excerpt 'How to design learning experiences for millennials', we can ask ourselves:

4 *Is the language vague? The answer to this is yes. What does the author mean by 'millennials'? What does 'learning differently' mean? Learning what and in what context? How is a millennial's 'lifelong relationship with technology' different from other people's often long relationship with technology? For example, we're [the authors] not millennials but we have a 'lifelong relationship with technology' ourselves. After all, there have been machinery and devices – even electronic – for a long, long time.*

5 *Is the language emotional?* Well, the article also states that 'Millennials are the future, so it's critical that organizations design learning strategies that ensure they get the most out of this unique generation.' This implies that something bad will happen if we don't adapt how we design learning experiences. Typical '**moral panic**' language. So, the wording plays with emotions without making things explicit.

6 *Is it 'hyped up'?* Yes, big time. It's trendy to talk about millennials (or other generational 'groups') as if they're a different species and it's also popular to pretend that we need to adapt the design of learning experiences to people's preferences (yes, the **learning styles** myth is still flourishing, unfortunately. We discuss this further in Part II).

Step 1b: Flip it

'Flip it' means that you try to turn the argument upside down. In this case, we're asked by Joe Doe to take action. We're asked to design learning experiences differently for millennials, because otherwise… otherwise what? Ask yourself what happens if you *don't* take this age group into consideration and you design just as you've always done, based on what people need to do their jobs.

Willingham recommends writing down the following statement:

> If I do X, there is Y per cent chance that Z will happen.

In this case, the percentages don't really work, and that confirms how vague the article is. For example, the author suggests that we should focus on video when designing learning experiences for millennials. The question then is, if we use videos as learning experiences for millennials, whether we'll have Y per cent chance that they'll achieve their objectives. This of course is very hard to determine, which is why we need to raise our eyebrows and be suspicious.

Step 2: Trace it

This comes down to: Don't just trust what people say because it sounds good or 'logical' or because they're seen as an authority or a (self-proclaimed?) expert or even the company that sells the intervention. This doesn't mean we have to extensively research everything, but we need to dig a bit deeper and ask ourselves what kind of evidence there actually is for the claim. What kind of resources has the article used? Just take a critical look. In the example, the research centre 'We Do Research' is mentioned, so the least we should do is trace those references and determine the quality. Has it been published? In what journal? Is the research centre dedicated to a certain philosophy or, worse, to a commercial product? In other words, is the company doing the research or paying for it the same as the company that makes the product or service being researched? Is there corroborating evidence from other peer-reviewed, scientific articles?

This brings us to step 3....

Step 3: Analyse it

This step requires some basic statistical knowledge, but a critical eye can bring us quite a long way. Willingham suggests that if something sounds too good to be true, then it probably is. Similarly, if a claim sounds very strong, too generic, too dramatic, then it probably needs more nuance at best or is pure nonsense at worst!

In the context of the sample, how likely is it that, first, video is the *only* medium that grabs the millennial's attention, and second, that a millennial's brain processes a video 60,000 (or whatever the number is) times faster than other generations? Plus, even if this were true, we need to wonder what that means for learning through video. It's definitely worth analysing all these types of statement in more depth.

We also recommend that you look for and find people who do high-quality research-to-practice work and then use their work, which usually includes practical examples, to your benefit. In our field, people such as:

- Will Thalheimer
- Patti Shank
- Julie Dirksen
- Connie Malamed

- Clark Quinn
- The Learning Scientists
- David Didau
- Daniel Willingham
- Blake Harvard

This is just to name a few who do a really good job, not to mention what we do ourselves. When you read them, they refer to research to support their claims and are very cautious about what they claim. This doesn't mean that you should believe *these* people blindly, but it will make it easier to trace and get a feel for the research that's out there so that we can use it to design effective, efficient and enjoyable learning experiences, based on the learning sciences.

Step 4: Should I do it?

In our profession, most of the time this would be about questions such as should we apply this method, implement this strategy, buy this tool and so forth. After we've stripped, flipped, traced and analysed it and the results are positive, then it might be worth a shot!

Willingham's steps are one way to take a critical look at something we're reading, seeing or hearing so that we can decide if it's true and/or worth using. Although we don't all have to dive deeper into the learning sciences, it's worthwhile to understand what they are and why they're so important for us as learning professionals. So, let's find out.

What is the learning sciences and why does it matter?

The learning sciences is an interdisciplinary field, focusing on progressing scientific understanding of learning. It also engages in the design and implementation of learning innovations, as well as the improvement of instructional methodologies. Research in the learning sciences traditionally focuses on cognitive-psychological, social-psychological, and cultural-psychological foundations of human learning, and on the design of learning environments. Major contributing fields include **cognitive science**, computer science, educational psychology, anthropology and applied linguistics. The learning sciences is based on research emerging from cognitive science, behavioural science, computer science, philosophy, sociology and information science (Figure 1.4).

FIGURE 1.4 The learning sciences model

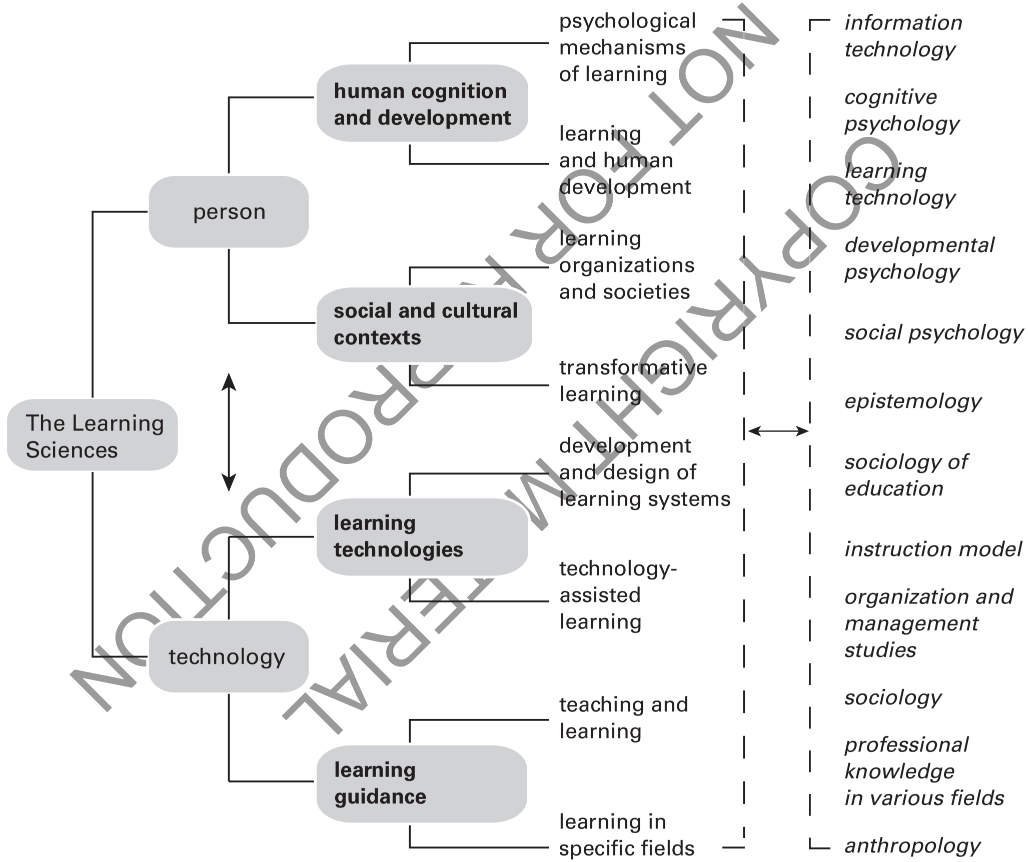

The learning sciences as a research field has only been around since the beginning of the 1970s. Keith Sawyer, a US psychologist who studies creativity, learning and collaboration, explains that before that, even when all major industrialized countries started to offer formal schooling to all children, there was no real knowledge of how people learn (Sawyer, 2005). Basically, schools were designed around assumptions (sometimes common sense, sometimes not) and there was no way to know what actually worked, what didn't work, and why. The good news is, now there's a way.

Though there's still a lot to discover, learning scientists have found common ground, after 50 years of research, on certain aspects of learning. For example, there's agreement that we need to:

- **Focus on deeper conceptual understanding.** To put it simply, factual and **procedural knowledge** is extremely useful and also necessary (how can you think deeply and/or conceptually about something that you know nothing or little bout?), but you also need to know when to apply it and how to adapt it to new situations.
- **Focus on learning *and* instruction.** Instruction alone, even high-quality instruction, isn't enough. We all need to actively participate in our own learning process. Without this, we won't be able to achieve deeper conceptual understanding of what we've learned.
- **Create learning environments.** The learning sciences has identified key features of learning environments that we need in order to acquire the full range of knowledge required for expert performance (facts, procedures and deeper conceptual understanding).
- **Build on prior knowledge.** No one's brain is a blank slate (**tabula rasa**). We all have acquired knowledge in one form or another, be it generally accepted 'truths' or conceptions about how the world works, some correct and some not (eg misconceptions). David Ausubel wrote in 1968: 'The most important single factor influencing learning is what the learner already knows. Ascertain this and teach him [*sic*] accordingly.' We learn best if we're in an environment that builds on, or at least acknowledges, our existing knowledge.
- **Include opportunities for reflection.** We learn better if we think about what we're learning. This can happen through pausing to think (with or without being prompted to do this), conversations with others, writing about it or creating some kind of artefact with it. In this way we have an opportunity to reflect on and analyse what we've learned. We all know

that, in the workplace, it's still generally accepted to let people participate in day-long training sessions with crammed agendas. Something to reflect on, we'd say.

- **Focus on learning transfer (performance), not just on comprehension.** This might seem obvious, especially in a workplace learning setting, but it's important to emphasize. Elizabeth and Robert Bjork (2011), among many others, have done extensive research not only around what people need to understand and remember (storage strength, reflecting how engrained or interconnected a memory representation is with related knowledge and skills), but also on being able to retrieve knowledge (retrieval strength – reflecting to what extent you can access the memory representation) and use it in the required context at hand.

Again, these basic facts come from 50 years of incremental research. And there's more. Many other learning principles that we know to be effective today are grounded in sometimes earlier influences on the learning sciences, such as **behaviourism**, or on constituent areas such as cognitive science, **constructivism**, **educational technology**, linguistics, socio-cultural studies and the information sciences. This is where the roots of the learning sciences lie. We don't necessarily need to dive deep into these theories to learn how to design effective, efficient and enjoyable learning experiences, but it's worth being aware of what they are.

> Though interdisciplinary, the learning sciences is most strongly grounded in psychology, and thus we've chosen to focus here on behaviourism, cognitive science and constructivism (Ertmer and Newby, 1993).

Behaviourism

Behaviourism sees learning as a change in observable performance. When looking at learning through a behaviourist lens, it occurs when an individual demonstrates a correct response or behaviour. The key elements are a stimulus, a response and an association between them. Behaviourism focuses on how the association between a stimulus and a response is made, strengthened and maintained. It's about the consequence of correct or incorrect performance. If correct, the performance needs to be reinforced in a positive manner, while, if incorrect, there will be negative reinforcement or punishment.

Behaviourism doesn't attempt to determine the structure of someone's knowledge nor if they're performing based on the correct mental processes.

EXAMPLE Behaviourism in today's workplace: Catch the Phish!

Compliance training, what's not to love… In this scenario, the organizational e-mail service includes a button that can flag 'phishing' or junk e-mail. Intermittently, people in the organization receive a dodgy e-mail. As soon as they identify an e-mail as being suspicious, they select the Phishing/Junk button. When they've identified an e-mail correctly as being 'risky' or 'non-compliant', they receive a message, praising them for being so diligent. If they fail to do so and respond to the e-mail or open a link that is included in the e-mail, they receive a message pointing out that they've done it wrong and that they've not met their responsibility to protect the organization's security. When the latter happens on a regular basis, they might get an official warning, or even worse. And there's a 'catch' as well, of course, when people get really afraid of clicking into an insecure e-mail or link and then they overcompensate by flagging real and trustworthy e-mails. What's missing in this process is the due diligence to a) understand if someone has identified the suspicious e-mail on the correct grounds (perhaps they think an e-mail is suspicious for the wrong reasons?) or b) understand what knowledge the individual lacks to help them correctly identify suspicious e-mails. To sum it up, there's a response to the performance, but not to the reasons behind that performance.

Cognitive science

Cognitive science is, in itself, an interdisciplinary study that researches the mind and its mental processes. It also considers what we know from fields such as sociology, linguistics, artificial intelligence, philosophy, **neuroscience**, psychology and anthropology about how people use knowledge in everyday settings (Thagard, 2008). Cognitive science focuses on things such as perception, memory, attention, reasoning and emotion. At first, the focus in this field was primarily on research in unnatural laboratory settings with strict methodologies. But somewhere in the 1990s, the influence of learning in more natural settings increased, and now cognitive science is very central to the learning sciences. Concepts such as expertise, reflection and problem-solving are all rooted in cognitive science and all have practical implications for how we design learning experiences.

Constructivism

Constructivism is actually a philosophy and not a psychological or instructional theory. According to constructivists, we ourselves impose meaning on the world. This means that there are multiple ways to structure and understand the world around us and that there's no such thing as an 'objective truth'. In other words, there are many meanings or perspectives for any event or concept and this is determined by our previous experiences. What one person has experienced is different from all others and they see, interpret and understand things differently from all others. One example is Brown, Collins and Duguid's (1989) situated cognition, meaning that knowledge is a product of the activity, context and culture in which it's developed. Knowledge is fundamentally 'situated' in the specific context that it's learned. Again, a critical component of constructivism is that there's no common, shared reality. Instead, reality is the result of constructive processes.

Of course, proponents of any of these influential theories would argue that 'their' theory is better than the others. But what's important to keep in mind is that, no matter how we use these theories, we always need to interpret them in the context that we're working in and only then we can determine their usefulness and meaning.

In addition to these influential theories, there are others that we definitely need to be aware of when designing learning experiences. First, we zoom in on neuroscience and **artificial intelligence (AI)** for at least two reasons. First, we need to understand what they are and how they help us design better learning experiences. And second, they are very hyped at the moment and thus we also need to know what they're not.

Neuroscience

Neuroscience is an interdisciplinary field, just like the learning sciences. The field works closely with disciplines such as mathematics, linguistics, engineering, computer science, chemistry, philosophy, psychology and medicine. Neuroscientists study the brain at a cellular, functional, behavioural, evolutionary, computational, molecular and medical level (Nordqvist, 2018). In the learning field, there's a lot of talk around how neuroscience can help us design better learning experiences. Thalheimer (2018) has explored what researchers such as John Medina (neuroscientist), Dan Willingham (psychologist) and Jeffrey Bowers (neuroscientist and cognitive psychologist) say about the influence of neuroscience on learning and instruction. When we summarize the findings, we must conclude that so far, the overall state of our knowledge on how the brain works is not yet very mature and, more specifically, that there are hardly any examples of novel and useful suggestions for learning and instruction based on neuroscience thus far. Of course, this doesn't mean that neuroscience won't offer any useful insights in the future, but thus far, we need to be cautious. In Chapter 6, Daniel Ansari, a professor in developmental cognitive neuroscience, gives us some interesting insights into how we should interpret findings from neuroscience and when and why we need to be cautious.

AI

Roger Schank (1987), a US artificial intelligence theorist, cognitive psychologist and learning scientist, explains that what AI is depends heavily on the goals of the researchers and that the definition of what AI is depends on what methods are used to build AI models. Roughly, so says Schank, AI has two main goals. The primary goal is to build an intelligent machine and the second is to discover the nature of intelligence. The main focus in the learning field, when it comes to AI, is on adaptive learning experiences.

For adaptive learning, what this usually refers to is something similar to the following (Paramythis and Loidl-Reisinger, 2003):

1. An algorithm that monitors the activities of the learner and interprets these on the basis of domain-specific models.
2. The algorithm infers learner requirements or preferences from the interpreted activities.
3. The algorithm acts upon the available knowledge on the learner and the subject matter to facilitate the learning process.

Remember here that an algorithm is neither objective nor infallible. As Cathy O'Neil (2017) tells us in her book *Weapons of Math Destruction: How big data increases inequality and threatens democracy*, an algorithm is created by people and what they plug into it – including all of their explicit and implicit preferences and prejudices – determines how it works and what it does.

An interesting question is: To what extent have we been able to achieve these goals in a learning context?

INTERVIEW Expert insights

Vincent Aleven on AI for learning

Vincent Aleven is a professor and director of undergraduate programmes in human–computer interaction at Carnegie Mellon University. His research focuses on learning technologies that use various types of AI to personalize instruction, such as (cognitive) intelligent tutoring systems and cognitive tutor authoring tools (CTAT).

Artificial Intelligence in Education (AIED) seems to be a pretty broad term. Can you please explain what your definition of AI in learning and education is?

Interestingly, AI is usually defined as a computer performing a task that traditionally can only be done by humans. But what we see is that actually AI sometimes performs tasks that humans can't necessarily do. A narrow definition of AIED could be 'educational applications of AI techniques' and a very broad definition could be something like 'investigating learning to support formal education and lifelong learning, using methods, tools, and representations of AI'. The field of AIED has radically evolved and broadened over the last 25 years or so, collaborating closely with communities such as computer-supported collaborative learning, the learning sciences, learning analytics and so on.

Recently, researchers and developers have started to focus on designing and implementing effective 'human–AI partnerships', which capitalize on complementary strengths of the AI and the humans involved (eg learners, instructors, administrators etc). In creating such partnerships, it's important to support and promote important human values such as competence and autonomy. For example, an AI-based system might detect that a learner is struggling but might not 'know' why. It might provide some help, perhaps a helping strategy or explanation it has not tried before, and it might in addition inform the instructor that something unusual is going on. The instructor, upon being notified by the AI, would understand that the given learner might be having a bad day due to problems at home. They might offer social-emotional support. That way, the instructor and AI can form a synergistic partnership.

AI can be used in a number of ways: to follow – together with learners – their learning activities, assess their progress and learning, provide a degree of support and personalization to learners, keep instructors aware and informed of what's going on, and/or help the instructor help learners, especially those who might need help beyond what the AI-based learning software can offer. AI, machine learning and data mining can also be used in offline analyses of learning, to advance the scientific understanding of learning but also to help learners improve the educational technology so it serves learners and instructors better.

Your work focuses on two main themes: intelligent (cognitive) tutor systems and adaptive learning technologies. Could you explain what they are?

Cognitive tutors are one form of AI-based educational computer programs. They're grounded in cognitive theory and constructed around computational cognitive models of the knowledge that they aim to help learners acquire. These cognitive models generally represent learner thinking or cognition in the domain at hand. They can also include a representation of learner strategies and of typical misconceptions. Cognitive tutors support learning-by-doing, as learners practise solving complex problems, with intense guidance from the system. For example, they provide rich graphical user interfaces to provide a workspace in which learners can solve problems. They use their cognitive model to follow the learner as they solve problems so as to track learners' performance and knowledge growth. They can also provide personalized and adaptive tutoring, including just-in-time help and feedback during the process of solving, as well as personalized problem selection. Many (though not all) scientific evaluation studies have found that cognitive tutors help improve learners' learning outcomes. The technology has been commercially successful and is being used in middle schools and high schools in the United States by roughly half a million learners yearly, primarily in mathematics courses. Cognitive tutors thus are an interesting example of successful transfer of cognitive science and learning sciences research to actual educational practice.

Cognitive tutors are one form of adaptive learning technologies. In general, learning technologies are adaptive to the extent that they take into account that learners differ in their needs and interests as well as that they change over time. Perhaps more surprisingly, adaptivity can also mean that the designers created or revised the system based on a deep data-informed understanding of how learners often have the same experiences within a given task domain. As my colleague Ken Koedinger says, in many domains, different learners quite often experience the same challenges, although some get there earlier, others later. Instruction designed to address common difficulties, or take advantage of common prior knowledge or strategies, may be called 'adaptive' as well.

What exactly does that personalization and adaptivity look like?

In any form of adaptivity, two key questions are 'What to adapt?' – namely, features of the instruction – and 'What to adapt *to*?' – namely, learner variables or characteristics. For both of these questions, there are many different options. For example, a system might adapt to a learner's evolving knowledge state or a learner's specific solution path through a problem. Especially when the problems that learners work on are complex, it's important that the system can recognize and accept different solution paths. Systems might also adapt to learners' affective state, aspects of their motivation, and even to how they regulate their own learning. They also might adapt to social factors, or to how learners collaborate. In principle, any feature of the instruction can be adapted in response to one or more of these learner variables. The system must have a way in which to assess or track these variables, so it can respond appropriately. This area has traditionally been called 'learner modelling'. Over the years, many techniques have been developed, and extensively studied, in the area of AIED but also in the areas of educational data mining (EDM) or learning analytics.

A well-known form of adaptivity is mastery learning, which is implemented, in one form or another, in many forms of learning software, including cognitive tutors. The tutoring system tracks each individual learner's knowledge, in terms of the detailed knowledge components, based on that learner's performance within the given sequence of the problem. Based on this assessment, learners receive individualized amounts of practice to master all the knowledge components. The tutor lets the learner get to the next unit only when they've mastered all the skills in the current unit. The learner does not have to take a test in order to 'convince' the system they have mastered the skills of a unit, because the system continuously assesses the learner – its learner model captures knowledge growth over time.

Other examples of adaptivity are systems that (based on the well-known **expertise reversal effect**) transition from having learners explain examples to having them solve open problems at just the moment where the examples start having diminishing returns, which may be different for each individual. Also, such systems might give feedback not only on task performance at the domain level but also on how learners self-regulate their own learning (eg seek help as needed).

What makes this type of system 'intelligent'?

A short answer is that a system is 'intelligent' to the extent that it adapts to differences and similarities between learners – everything we talked about before. In this context, I view the terms 'adaptive' and 'intelligent' as rough synonyms.

An interesting implication is that it's the behaviour of the system that determines whether we view the system as intelligent, not so much that there's a particular kind of algorithm or knowledge representation under the hood. This is one reason that I haven't singled out any particular AI technology. Though of course the choice of AI methods or algorithms matters when you come to a third key question, 'How to adapt?'

How is AIED currently applied in organizations/workplaces and how could they potentially be applied?

In the workplace, you see examples like language learning or many tasks in a military context, such as equipment maintenance, electronics troubleshooting, medical first-aid procedures or even team training. Lewis Johnson (for example, Johnson and Valente, 2009) has done very interesting work on using virtual reality for language learning as well as cultural learning. Another example of AIED in the workplace is iHelp, an older learning content management system, that started as a peer help system but is now more like an ecological approach. It's a bit complicated, but basically it mines all kinds of data around what choices learners make to achieve a learning goal and how successful they are with that. The system looks for patterns and can then, for example, recommend the best next step. Or recommend a peer to contact when the learner ends up in an impasse during the learning process.

Intelligent tutors or adaptive systems work particularly well when you have clear learning goals. But in the workplace, things might change too fast to make it worth the investment. These systems are quite expensive to design and develop. So, unless you have skills that a large number of workers need to learn, it wouldn't necessarily be the most obvious choice. Recent developments in tools for rapid authoring of systems are worth keeping an eye on, as they might lower the threshold for tutor development.

Are there examples of cognitive tutors for ill-structured problems? That could be very useful in the workplace.

Cognitive tutors are often applied in STEM (science, technology, engineering and mathematics) domains, to well-structured problems. We don't have many applications of cognitive tutors supporting ill-structured problems, with a few exceptions in domains such as language learning, intercultural competence and causal argumentation. This isn't necessarily a limitation of the technology itself. The technology can handle problems that have a wide variety of solutions. Also, cognitive task analysis can oftentimes bring helpful structure to domains that previously were unstructured. But a key challenge is, if you want to support more open-ended problems, to provide an interface that allows for many different solution approaches without prompting any specific ones. Also, it's challenging to identify strategies for helping learners with open-ended problems that honour – don't trample on – the essential open-ended nature of the problem. It's even more challenging to create AI-based systems in domains where even experts disagree, such as for example in the legal domain. Systems in such domains need to implement accepted methods of reasoning or analysis, as well as teaching. Modelling those can require pushing the envelope in use of AI methods. As another example, it would be challenging indeed to create cognitive tutors that support project-based learning, although they could likely be helpful by providing the just-in-time learning that often needs to happen in that context and that can be challenging for instructors to orchestrate.

Is there a gap between what we expect from/hope for in AI and what it can actually deliver?

I tend to think of it as human–AI partnerships. I think, in this particular context, that AI can definitely do things that instructors can't, because they might have a large group of learners to support, and they can't give each learner the right amount of individualized attention. These tutoring systems help with that. This is an example of what I said earlier around AI doing things that we, as humans, are actually *not* capable of.

On the other hand, there might be things that AI could help with, but I think in the end, what learners really care about is if the instructor cares about them, understands and empathizes with their struggles, celebrates their successes with them. The instructor also has a much wider repertoire of instruction strategies than the system.

There are examples of AI drawing attention to learners. As an example, from research done by PhD candidate Ken Holstein, instructors would wear mixed reality smart glasses and the intelligent tutor presents them with real-time analytics to draw their attention to, for example, learners who aren't active, are gaming the system, make lots of errors or avoid hints when they're likely needed. This is a great example of a human–AI partnership as the technology, designed very carefully with lots of input from instructors, helps the instructor to tune in to the learner who most needs it at that point in time. In a classroom study, when instructors had the glasses, learners learned better!

I see many ways in which the fields of AI and human–computer interaction (HCI) – and both are critical – can jointly make progress and tackle ever more challenging educational and learning problems, through better language understanding, better, more wide-reaching pedagogies rooted in deeper learner modelling methods, with better design not just of the user interface, but of the fundamental role that technology plays as it partners with humans.

The last science we like to dive a bit deeper into is educational technology. One stubborn myth is that new technology is causing a revolution in the space of instruction and learning (de Bruyckere, Kirschner and Hulshof, 2015). We feel that it's important to discuss educational technology separately because technology in general plays a huge role in most people's day-to-day lives and the least we should do is understand how technology should or could be used to support learning effectively.

Educational technology

The Association for Educational Communications and Technology (AECT), the professional association of instructional designers, educators and professionals, defines educational technology as 'the study and ethical practice of facilitating learning and improving performance by creating, using and managing appropriate technological processes and resources' (Richey, 2008: 24). Note that technological processes go a lot farther and are a lot older than information and communication technologies. When the AECT began in 1923 (!) it was founded as the National Education Association's Department of Visual Instruction.

EDUCATIONAL TECHNOLOGY AND EDUCATIONAL REVOLUTION

It's interesting to note that it was in the same year as the 'founding' of the AECT, 1923, that Thomas Edison repeated an earlier prediction, so tells us the Associated Press on 15 May, that motion pictures would replace books as the primary instructional medium; he figured it would take about 20 years to complete the conversion. This claim has been made for just about all of the technological 'revolutions' since then!

Since the 1990s there has been a huge push to implement computers in learning environments and the investments have been huge (Halverson and Smith, 2009). Unfortunately, these investments have not really paid off. There's a simple reason for that.

Ongoing research shows that it's not the technology in itself having a positive impact on learning. Based on his research, Richard Clark concluded in his provocative 1983 article that media are mere vehicles that deliver instruction but do not themselves influence learner achievement. He compares media with the truck that delivers our groceries. The truck is only a vehicle and does not in any way cause changes in our nutrition. He suggests focusing on effective instruction and learning strategies instead. Clark also recommends exploring how difficult or interesting learners find various media. This might help to find relations between media and the willingness of learners to invest in learning.

Later work by Clark and Feldon (2014) in their chapter 'Ten common but questionable principles of multimedia learning' confirms that the effectiveness of learning is determined primarily by the way the medium is used and by the quality of the instruction accompanying that use. When technology is used to support instruction or learning, the choice of the technology is not what has an impact on the effectiveness of learning.

Not to be overlooked is the fact that educational technology isn't new. Actually, both the book and the blackboard (the slate thing that was replaced by the greenboard and the whiteboard and not the content management system) are examples of educational technology. The first educational software goes back as far as the 1960s and was based on Skinner's behaviourist theory. Basically, since Edison's prediction, it's been same old, same old since 1923.

CASE STUDY

Skinner's teaching machine

Skinner's (1958) device (Figure 1.5) might look nothing like some of today's fancy-looking e-learning programs; the underlying reasoning is still pretty much the same. The device enables people to drive their own learning. His teaching machine shows clear similarities with current 'fill in the blank' computer-generated questions. This type of question was, according to Skinner, for 'more advanced learners'. He explained that such questions are more effective than multiple choice, because learners have to compose their own response rather than select it from various options. This means that the focus is on recalling, instead of recognizing. Another interesting consideration that Skinner made was that, when someone has to acquire complex behaviour, they must go through a carefully designed learning process. Skinner explained that each step of the required sequence must be small enough to be taken successfully and it must help the learner to move closer to the fully competent behaviour. Hence, it's the job of the teaching machine to ensure that the required steps are taken in a prescribed order (see Table 1.1).

FIGURE 1.5 Skinner's teaching machine

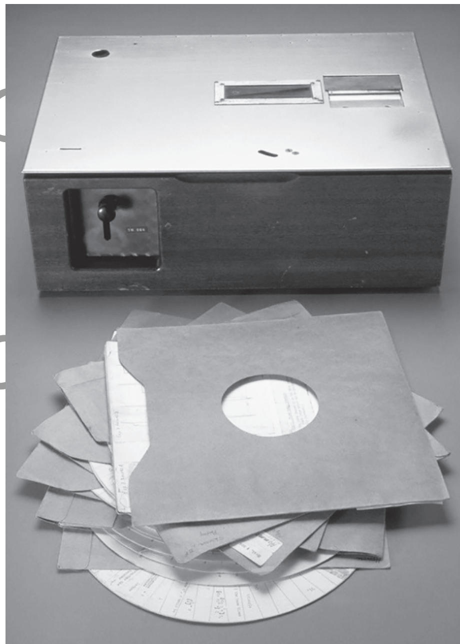

Reproduced with kind permission from the Smithsonian National Museum of American History

TABLE 1.1 Example of recalling activity of Skinner's teaching machine

Sentence to be completed	Word to be supplied
When the hot wire glows brightly, we say that it gives off or sends out heat and __________.	light
The higher the temperature of the filament, the ______ the light emitted by it.	brighter, stronger

The material of Skinner's teaching machine is printed in 30 radial frames on a 12-inch disk. The learner inserts the disk and closes the machine. All but a corner of one frame is visible through a window. The learner writes his or her response on a paper strip, which is exposed through a second opening. When ready, the learner lifts a lever on the front of the machine and uncovers the correct response in the remaining corner of the frame. If the two responses are the same, they move the lever horizontally. This movement punches a hole in the paper opposite the response. This way, the correct answer is recorded. As a consequence, the machine is altered and won't show the same frame again. No matter if the response was correct or incorrect, a second frame appears when the lever is returned to its starting position. The learner continues until all questions have been answered. Next follows another round, in which only those questions that were initially incorrectly answered appear.

In summary, the machine does a couple of things. First, it increases efficiency, as no instructor is required to teach these materials. Second, there's immediate feedback after learners provide an answer, hence they have an opportunity to continuously reflect on their own progress. The machine also identifies the weaknesses of the learner. After all, it will repeat the questions that have been answered incorrectly. Furthermore, the machine provides a personalized experience as the machine only presents that material for which the learner is ready. It asks the learner to take only that step that they're ready to take. Our conclusion: not bad for the 1960s, but not very different from many of the computer- or tablet-based 'modern' drill-and-practice things we use today.

Let's recap why we, as learning professionals, need to know about the learning sciences and its history. First, let's envisage what happens when we design learning experiences *not* based on the learning sciences. In essence, this means we're regressing to the early 20th century, when we still designed formal education based on beliefs and common sense. Unfortunately, years of experience doesn't mean you're making informed decisions and doing it right. Doing things over time gives you experience, yes, but that doesn't mean that you know what it's worth. There was a lot of experience with trephining and bloodletting before we learned that neither technique worked. The decisions we make based on experience might be well informed or they might suffer from knowledge gaps, the misconceptions we've developed over time, the **Dunning–Kruger effect** (this means that incompetent people are not capable – *because of* their incompetency – to see that their reasoning, choices and/or conclusions are just plain wrong; also see Chapter 11), and so forth.

We need the learning sciences because we need to make evidence-informed design decisions based not on what we *believe* but rather on what we *know*.

Knowledge grounded in science increases over time and progresses incrementally. Of course, it's an ongoing learning process in itself. But that's not a reason not to leverage what we know works at this point in time. Or to avoid what we know is ineffective.

Some people hide behind the fact that evidence from learning sciences is not always very strong. And although this can be true (as explained earlier, this is one reason why we use the term 'evidence-informed'), this is a wrong reason to not use the evidence that we have. If the evidence is weaker, we just need to be more cautious and aware of contextual variety. Modern medicine, biology, physics; science is the only reason these fields have progressed to where they are today.

Let's be humble and acknowledge that many things we talk about today as 'being innovative' are as old as you, your mother or grandmother (depending on your age). Think microlearning, animation, digital, AI, and the just discussed teaching machine. Let's stop wasting time, effort and money on reinventing wheels that might not need to be reinvented because they might already have led to dead ends at some point in time and let's start spending time on leveraging the evidence we have to design learning experiences in an evidence-informed way. This brings us to the next chapter in which we discuss what learning experiences are, what it means to design them and how to design what we call 3-star (effective, efficient and enjoyable) learning experiences.

Chapter 1 key points

1. Approaching learning experience design in an evidence-informed way means that what we do is backed by research from the learning sciences. Only when we use the evidence available to us as practitioners can we improve our practice through making well-informed decisions. It prevents us from wasting time and money and doing potential harm to learners. It also increases our credibility as a profession.
2. The learning sciences gives us insights into how to design effective, efficient and enjoyable learning experiences. It's an interdisciplinary field, focusing on progressing scientific understanding of learning, the design and implementation of learning innovations and the improvement of instructional methodologies. Though there's still a lot to discover, there's definitely common ground on certain aspects of learning, such as the importance of prior knowledge and **learning transfer**.

3 We can use Daniel T Willingham's steps to our benefit when trying to judge a piece of information. The steps are: 1) strip it and flip it, 2) trace it, 3) analyse it and 4) should I do it?

References

Ausubel, D P (1968) *Educational Psychology: A cognitive view*, Holt, Rinehart and Winston, New York

Bjork, E L and Bjork, R A (2011) Making things hard on yourself, but in a good way: creating desirable difficulties to enhance learning, in *Psychology and the Real World: Essays illustrating fundamental contributions to society*, ed M A Gernsbacher *et al*, pp 59–68, Worth, New York

Brown, J S, Collins, A and Duguid, P (1989) Situated cognition and the culture of learning, *Educational Researcher*, **18** (1), pp 32–42

Clark, R E (1983) Reconsidering research on learning from media, *Review of Educational Research*, **53**, pp 445–59

Clark, R E and Feldon, D F (2014) Ten common but questionable principles of multimedia learning, in *The Cambridge Handbook of Multimedia Learning*, ed R Mayer, pp 151–73, MPT, Lancaster

de Bruyckere, P, Kirschner, P A and Hulshof, C D (2015) *Urban Myths About Learning and Education*, Academic Press, San Diego, CA

Ertmer, P A and Newby, T J (1993) Behaviourism, cognitivism, constructivism: comparing critical features from an instructional design perspective, *Performance Improvement Quarterly*, **6** (4), pp 50–72

Halverson, R and Smith, A (2009) How new technologies have (and have not) changed teaching and learning in schools, *Journal of Computing in Teacher Education*, **26** (2), pp 49–54

Johnson, W L and Valente, A (2009) Tactical language and culture training systems: using AI to teach foreign languages and cultures, *AI Magazine*, **30** (2), pp 72–72

Nordqvist, C (2018) [accessed 20 September 2019] What is neuroscience?, *Medical News Today*, 26 June www.medicalnewstoday.com/articles/248680.php (archived at https://perma.cc/XYZ8-9Z2E)

O'Neil, C (2017) *Weapons of Math Destruction: How big data increases inequality and threatens democracy*, Broadway Books, New York

Paramythis, A and Loidl-Reisinger, S (2003) Adaptive learning environments and e-learning standards, In *Second European Conference on E-Learning*, vol 1, pp 369–79

Richey, R C (2008) Reflections on the 2008 AECT Definitions of the Field, *TechTrends*, **52** (1), pp 24–25

Sackett, D L *et al* (1996) Evidence based medicine: what it is and what it isn't, *Clinical Orthopaedics and Related Research*, **455**, pp 3–5

Sawyer, R K (ed) (2005) *The Cambridge Handbook of the Learning Sciences*, Cambridge University Press, Cambridge

Schank, R C (1987) What is AI, anyway? *AI Magazine*, **8** (4), pp 59–65

Skinner, B F (1958) Teaching machines, *Science*, **128** (3330), pp 969–77

Thagard, P (2008) *Hot Thought: Mechanisms and applications of emotional cognition*, MIT Press, Cambridge, MA

Thalheimer, W (2018) [accessed 20 September 2019] Brain Based Learning and Neuroscience – What the Research Says!, *Work-Learning Research*, 29 March www.worklearning.com/2016/01/05/brain-based-learning-and-neuroscience-what-the-research-says/ (archived at https://perma.cc/SN26-DV9C)

The Associated Press (1923) Edison predicts film will replace teacher, *Highland Recorder*, 15 May, p 2

Willingham, D T (2012a) Measured approach or magical elixir? How to tell good science from bad, *American Educator*, **36** (3), pp 4–12

Willingham, D T (2012b) *When Can You Trust the Experts? How to tell good science from bad in education*, John Wiley & Sons, Hoboken, NJ