CHAPTER 1

WHY COGNITIVE LOAD THEORY?

- **Students can be active, solving problems or completing investigations, without learning key concepts.**
- **Cognitive load theory attempts to explain why some instructional approaches lead to more learning than others.**
- **Cognitive load theory uses the scientific method and its insights are built on evidence from experiments.**
- **The *Worked Example Effect* was drawn from experimental evidence and demonstrates that novices learn more from studying worked examples than by solving equivalent problems.**

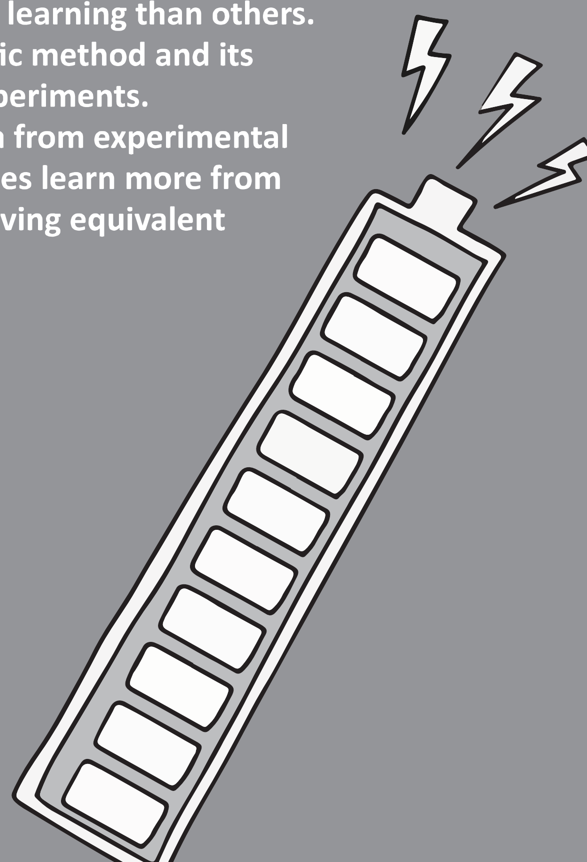

In 2017, Dylan Wiliam took to Twitter to write, 'I've come to the conclusion that Sweller's Cognitive Load Theory is the single most important thing for teachers to know' (Wiliam, 2017). Why would Wiliam, whose influential career has seen him advocate for formative assessment in the classroom and become Deputy Director of London's Institute of Education, come to such a conclusion? In this chapter, you will begin to find out.

In essence, cognitive load theory is simple – so simple, in fact, that it can seem like common sense. However, as anyone who has worked in schools or faculties of education can attest, its implications are routinely ignored.

Let us start with a finding that does seem genuinely surprising and is from one of the early, seminal experiments that John Sweller conducted with his colleagues, prior to founding the field of cognitive load theory (Sweller, Mawer and Howe, 1982). In the early 1980s, undergraduate students at the University of New South Wales were asked to solve a series of weird maths problems. They were told to proceed from a start number to a target number using a combination of only two rules – multiply by three or subtract 69. However, in the first four problems, the solution involved alternating the two rules, as in Figure 1.1.

Problem	Start number	Solution	Target number
1	60	×3, –69	111
2	81	×3, –69, ×3, –69	453
3	34	×3, –69, ×3, –69, ×3, –69	21
4	35	×3, –69, ×3, –69, ×3, –69, ×3, –69, ×3, –69, 3, –69	156

Figure 1.1 Problems students could solve without noticing the underlying pattern

Adapted from Sweller, Mawer and Howe, 1982

Some of the problem solvers were randomly allocated to a group that asked them to write out their answers and look for a pattern, whereas others were allocated to a group that were not told this. Although this latter group could generally successfully solve the problem, they rarely noticed the pattern, meaning they took much longer than the first group to solve the fourth problem.

Sweller and his colleagues conjectured that without a rule to follow, test problem solvers were forced to resort to 'means-ends analysis'. When evaluating each move, they judged whether it would bring them closer to the target or not. The process of means-ends analysis is taxing, and it occurred to Sweller that it was so taxing that problem solvers had no mental resources left over to notice patterns that were right in front of them.

The problem was that the test problem solvers were busy. They were actively and successfully solving problems, but they were not learning the structure behind those problems and so they could not transfer that understanding to new problems.

This experiment reminds me of my early adventures in science teaching and one practical activity I would organise for my students. It involved placing marble chips in hydrochloric acid and measuring the rate at which carbon dioxide gas was released. The key point was that the marble chips came in different sizes and therefore had different surface areas. By monitoring the rate of reaction, students were intended to conclude that the smaller the chips, the larger the surface area and the faster the reaction.

Students were exceptionally active during this experiment. They asked me lots of questions and I would give plenty of advice on how to carry out the procedure. To reluctantly deploy a word rendered almost meaningless by overuse, my students were *engaged*. It was a fun practical activity and unlike some of the duller experiments, they did not drift off into chatting to their friends about irrelevant matters – something I had already discovered was a feature of group work.

Yet at the end of all this, when the equipment was packed away on the trolley and the desks were wiped clean, few students knew the relationship between chip size and reaction rate and so I found myself patiently explaining what the students should have discovered.

In this case, my students were not using means-ends analysis – they were following a detailed set of instructions – but perhaps a similar phenomenon explained the lack of learning. The process of completing this experiment and executing the instructions was so taxing that mental capacity had effectively been exhausted and there was nothing left for noticing patterns.

As uncomfortable as it is, we need to come to terms with the idea that activity – even successful activity – does not necessarily lead to learning. Even when it does, it may not lead to the learning we intended, assuming we intended anything specific at all.

And this raises a question of definitions. One of the easiest ways to become insufferable on Twitter is to ask for definitions of commonplace words such as 'learning'. For now, I am going to park that issue and ask you to assign the term whatever meaning you wish. However, we will need to return to this, and we will see that in cognitive load theory, learning is defined in a specific way that prevents us being woolly in its application. Yet, being woolly about learning is the key feature of an approach that I label 'activity-based planning'.

A few years back in a school where I worked, there was a fashion for a mathematics game called 'Greedy Pig'. Greedy Pig is a game of probability where students are given the choice to either bank the points they have earnt or gamble them on the next roll of a dice (see Polster and Ross, 2014, for a review). A paraphrasing of the discussion between the teachers about Greedy Pig would be something like:

> 'Did you do Greedy Pig yesterday?'
> 'Yes. The kids loved it.'
> 'Oh, I'd better do it today then, so my kids don't miss out.'

At no point do I recall a discussion of what students were supposed to *learn* from playing Greedy Pig. Yet, activity-based planning of this kind is something I have engaged in frequently throughout my career and experience suggests it is widespread in schools.

Even when we have a clear and explicit rule – such as that a problem can be solved by alternating two moves; or the larger the surface area, the faster a reaction – we cannot be sure that it will be learned by students completing highly relevant tasks, so what are our expectations of activity-based planning?

Given its prevalence, there is the possibility that many students in many schools are often busy without learning much.

So, from the time of its gestation, cognitive load theory offers us an insight – that activity does not necessarily imply learning. Borrowing from the field of thermodynamics, we could perhaps posit this as the ‘zeroth rule’ of cognitive load theory. It is an insight that is perhaps obvious and yet its implications are routinely, daily violated. If cognitive load theory offered us this insight and nothing more, it would be an important thing for teachers to know. However, once we start thinking about how much we are taxing mental resources in the tasks we select for students to complete – the ‘cognitive load’ that gives the theory its name – there are plenty more implications for the way we teach.

We will return to those implications in later chapters, but first we must deal with the small matter of the truth and how we establish it.

Different fields have their own ‘epistemology’. This means they attempt to establish truth in different ways. Lawyers pursue truth by testing evidence against laws and precedent. Historians pursue truth by analysing historical sources. Scientists pursue truth through application of the scientific method – they propose a model, test it through observations and experiments, reject it if the results are at odds with the model and provisionally accept it if they are not.

A scientific theory – such as cognitive load theory – is one of these provisional models. It makes useful predictions in a limited range of situations, but we can never rule out that it will be superseded by a more complete theory. And in most cases, scientific models represent simplifications of the world and do not attempt to describe every detail of it fully and accurately – a point we will return to.

Unfortunately for scientists, ‘theory’ has an additional meaning in everyday language, that of ‘what someone thinks may be the case’. Such a theory does not have to pass the test of observations and experiments. There is no requirement for it to make accurate predictions. This kind of theory is a drag on education. Figures such as Rousseau, Spencer, Dewey, Piaget, Freire, and

Foucault have strongly influenced education, and yet any requirements that their notions pass scientific tests are loosely, if ever, applied.

The tradition of vesting authority in some prominent philosopher is so strong that there are those in education who view educational theory and evidence-based practice as being in opposition (Atkinson, 2000). They complain about calls to test ideas against evidence – pleading that the field of education is special and different to other fields that have adopted evidence-based approaches – and they describe attempts to apply evidence to education as 'positivistic' (Biesta, 2007). Such reactions are best understood as an immune response, given that a few relatively simple experiments could potentially prove much of the field wrong and so invalidate the life's work of most of its practitioners.

Which leads us to the series of relatively simple experiments published by John Sweller and Graham Cooper in 1985 that are central to the development of cognitive load theory.

Through a series of pilot studies, Sweller and Cooper carefully selected Sydney high school students from Years 8 and 9[1] who had relevant knowledge of basic algebra principles but who were not yet proficient at solving algebra problems that used these principles.

Each student met individually with an experimenter. No matter which of the two experimental groups they were randomly allocated to, the experimenter talked them through a couple of worked examples that demonstrated the use of the algebra principles involved in the study, took questions, and ensured the students understood the task.

One group of students then solved a series of paired problems that used these principles. Each pair had the same structure and required the same solution method. The other group were given the same pairs of problems but the first problem in each pair was a worked example that the student could study, like the worked examples presented at the start.

[1]Year 9 in Australia is equivalent to the Ninth Grade in the U.S. and Year 10 in the U.K.

Finally, all students were asked to solve the same series of test problems, which were similar to the ones they had just been solving.

In general, students who were given worked examples alongside problems to solve were quicker and more accurate at completing test problems than the ones who solved pairs of problems. It is a striking effect that has now been replicated many times in a wide range of different learning contexts (Sweller, Ayres and Kalyuga, 2011). However, there is an aspect of this original experiment that I think is often missed. The students who were given pairs of problems to solve were still given a couple of worked examples at the outset and the experimenter checked they knew the relevant algebra principles and understood the task. This was not a study that tested whether full guidance was more effective than no guidance *at all*, but a study that tested whether more guidance was better than less guidance. That distinction will become important.

This paper and the series of experiments it describes provided the first example within the field of cognitive load theory of what is now known as the *Worked Example Effect* – the first of several such 'effects' we will encounter. However, before we meet more effects, we need to look at how cognitive load theory models the mind.

IDEAS FOR THE CLASSROOM

It is worth considering the way you use worked examples in your classroom. Worked examples can include anything from traditional mathematics problems to modeling how to write a paragraph or explain a concept. When you demonstrate a worked example in class, do you pair it with an equivalent problem for students to solve? If so, is the problem highly similar to the worked example or does it draw on different knowledge and skills? When students are learning something new, the structure of the example and problem should be the same.

(Continued)

It can often be challenging to think of ways to break down complex tasks and create pairs of similar problems. Imagine, for example, a group of English teachers working together who have identified that students struggle with identifying the contention in a persuasive text. Modeling this process can seem challenging, particularly since such texts can be long. However, we can strip some of this back initially by using short arguments such as a letter to the editor of a newspaper. The teacher can line two letters up against each other, modeling finding the contention in the first argument, and then ask the students to find the contention in the second argument. Over time, the teacher can gradually build up to longer pieces of writing.

REFLECTION

1 **Think of a time your students successfully completed a task but you were uncertain that they had learned what you had intended. What was the task and what were its key features?**

2 **Has anyone ever suggested to you that something is a great activity without suggesting what students will learn as a result?**

3 **Is education a field like medicine that can be subjected to the scientific method, or is it a fundamentally different endeavour?**

4 **How do you use worked examples in your teaching? What do they look like? Do you present them at the start of each phase of teaching or thread them throughout?**

NOTE IT DOWN

Think of a favourite activity that you have completed with your students in class.

1 Briefly describe the activity

2 What did you intend students to learn as a result of this activity? Be as specific as you can.