3

LONG-TERM MEMORY: THE BASICS

INTRODUCTION

In the first chapter of this book, we introduced you to some of the main features of long-term memory (LTM). We will now explore how LTM works and the theory behind it, and clarify certain key principles that can be applied to all aspects of learning across the curriculum. Most importantly of all, we will explore how students can store things in LTM memory in way which will be resilient, well organised and retrievable.

A key question that you might have about LTM is around how long it lasts. How long is 'long'? We know that working memory is a very brief memory store; LTM comes into play for any information which we want to store beyond the current task. Even storing something for a few minutes while you work on another task would be considered a function of long-term rather than working memory.

LTM can therefore be seen as a permanent storage system, or at least one that *supports* permanent storage. How well things are actually accessible later on depends on a number of important factors which are important for teachers to understand. These will be explored in this chapter.

It is worth bearing in mind that LTM will be at work even when it is not being used effectively – for example, when students cram for tests and exams. These students are using LTM; it's just that they are not using it effectively. Perhaps rather than focusing on its 'long' duration, it is more useful to see LTM as the memory store which holds information while you are doing something else. If your students can retain something without paying attention to it – for example, while distracted by a conversation, or overnight – then we can say that LTM is at work.

All the same, we need to be aware LTM is subject to forgetting. It is a misconception to see our memories as a permanent recording and, of course, students who use flawed study strategies often cannot recall the bulk of what they have tried to memorise. You could consider the task of retrieving memories as being rather like finding something in a messy cupboard in your home. It might be in there, but you can't necessarily locate it when needed.

New knowledge and skills need to be stored permanently, but they also need to be *accessible*, so learners can retrieve them when needed. This includes applying what they have learnt in a new context or unfamiliar situation, when undertaking further study, or in the workplace.

THE BASIC OUTLINE OF LTM

Trying to sketch a basic outline of what LTM looks like in practice has proved even harder for researchers to do than is the case with working memory. There are multiple theories, and each focuses on particular aspects without necessarily capturing LTM as a whole system.

Perhaps the most widely used model is the idea that LTM is a simple storage system or 'box', to which working memory is connected and acts as a gateway. Together with the simple model of working memory shown in Chapter 2, this forms Atkinson and Shiffrin's (1968) 'multi-store' model of memory.

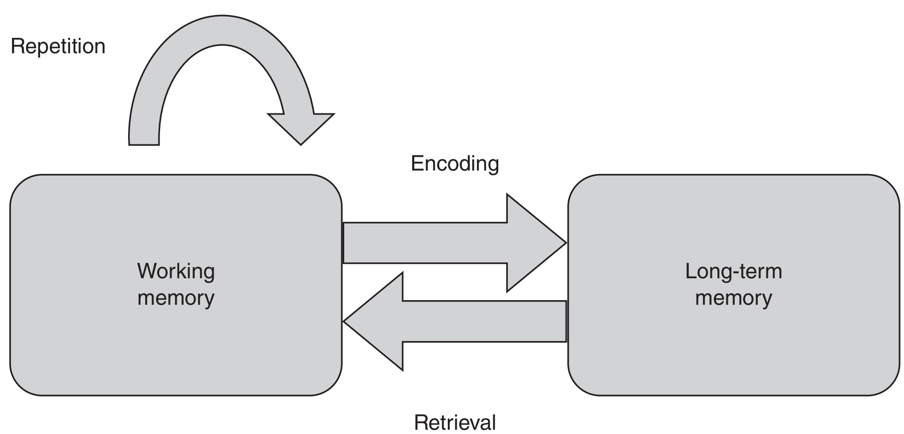

Figure 3.1 Multi-store model

The multi-store model suggests that WM and LTM connect together, and that they each work in fundamentally different ways. This idea is accepted by most memory researchers.

A key claim made by this model is that information needs to be processed in working memory in order to enter LTM, and the extent and duration of this processing correlates with how efficiently encoding happens. This idea is reflected in the following quote: 'Memory is the residue of thought' (Willingham, 2009). That is to say, the more we think about and process something in working memory, the more likely it is to be encoded to LTM.

The model does not illuminate what memories actually *do* once they are in LTM, or account for a memory being modified or edited after it is stored, something that happens often in education when learners are asked to reconsider or elaborate on information. It also does not tell us anything about how memories connect together. As discussed previously, memories cluster into networks called schemas (see Chapter 1), and this is important when students make connections across different school subjects.

Some of the main things this model fails to account for include:

- the fact that practice alone doesn't ensure memorisation;
- the role of active learning;
- the role of schemas, or previous learning already held in the LTM;
- the fact the content of LTM (or the knowledge and skills) influences processing in working memory;
- the fact some items (for example, stimuli which are dangerous or are personal to the learner) are more memorable than others;
- storage of different types of information (for example, verbal, visual).

These are significant omissions! As educators, we cannot ignore such major factors, as they affect whether curriculum-relevant material will be remembered or forgotten by our students.

The multi-store model is what theorists would describe as an information-processing model; LTM is compared to the storage of a computer. Many of the factors in the bullet points above don't apply to computer storage, but they *do* apply to real-world human learning. If anything, this just serves to remind us that computer analogies are limited. Like any analogy, this one has the potential to illuminate, but also the potential to mislead.

ACTIVITY

Try the following activity. Look at the following list of words for a ten to twenty seconds, then cover the page and write down everything you can remember.

Sharp	Syringe
Injection	Thread
Haystack	Knitting
Straw	Sewing.

The task you just tried is known as the Deese–Roediger–McDermott paradigm. Did you include the word 'needle' in your response? It's not there in the list! A great many experimental participants tend to include it. This demonstration shows the contents of LTM – in this case, the meaningful connections between words – can affect performance on a WM task.

In the classroom, the same phenomenon could occur when reading a story in English. As they read, the learners are making sense of the story using associations with existing knowledge to help them follow and understand it. The significance of the hand washing in the plot of Macbeth, for example, requires each learner to access schemas around guilt and contamination. This shows that long-term memories don't just sit passively in a box while students rehearse new things in working memory. Rather, memories form a network, and previously learnt information in LTM affects how learners understand incoming material.

THE ROLE OF FORGETTING AND SPACING

We will explore the multi-store model and address the 'gaps' in the model one at a time in order to get a fuller picture of how LTM operates in the classroom. However, first let's explore the role of spacing and decay of memories over time, which the model partially addresses.

LTM is subject to forgetting and this proceeds according to the forgetting curve (see Chapter 1), so students forget their learning fast at first, and later slower.

We call this forgetting process in LTM *decay*, meaning as time passes things start to disappear.

The concept of decay is included within the multi-store model, and fits with the need for consolidation of learning over time – something every teacher recognises as critical. As memories are forgotten, the model implies re-learning after a delay would be more beneficial than doing it straight away.

SCENARIO 3.1

Mrs T teaches her class about how to work out coordinates on a map in her geography lesson. She then sets an activity a few lessons later asking students to come in and start the task on their desk. The task is to identify and mark coordinates on a worksheet using the learning they did in their previous lesson. Mrs T finds some of the students have forgotten how to use coordinates. This skill needs to be refreshed in order for them to progress with the task.

The idea of delayed practice fits with the spacing effect, in other words, that learning is better remembered if practice is delayed. As Bjork and Bjork (2011) explain, it would be better for a learner to wait until they were on the verge of forgetting something before practising it again. Practice which is scheduled to happen after some forgetting has taken place – as in the example with Mrs T and her lesson on map coordinates in Scenario 3.1 – helps to secure new learning. And as time goes on, forgetting slows and retrieving the information again becomes easier and more automatic.

In the classroom, learners can be reminded of ideas just as those ideas are beginning to fade, and may remember things partially, or with errors. However, the multi-store model sees information as being 100 per cent present or absent in LTM, and has no way of explaining things being *weakly* present, or *partially* present. The modal model also doesn't provide any useful guidance on the timing of classroom practice at all.

Perhaps most importantly, we shouldn't see learning new material as being a one-off event. Instead, delayed consolidation is necessary to gradually boost the storage strength of a memory (see also 'A Developed Theory', below).

When making use of the spacing effect:

- learners will master something via an initial study session – in Mrs T's lesson above this was learning how to add coordinates to a map;
- learners will return to the material and practise it actively after a delay – in Mrs T's lesson this was a starter activity a few weeks after the initial learning;
- learners will revisit the material again on at least one occasion, with the practice now becoming more complex and interlinked with other material. In Mrs T's lesson students may learn to change incorrect coordinates on a map. With further practice and through linking in other knowledge, they will be able to explain the impact of inaccurate coordinates – for example, when the emergency services are attempting to locate someone for mountain rescue.

REFLECTION POINT

It is widely considered the spacing effect should be applied more widely in schools. How exactly would this work in your context?

ACTIVITY

Try to think of a specific way you could change the timing of practice for learners in your context. For example, is there a consolidation task or test which could take place after a delay of at least a week? Could a project task encompass meaningful review and consolidation of a previous topic and be part of the planned lesson regularly?

Remember, using the spacing effect well means learners should get a chance to study something in enough depth to grasp it on the first occasion. Then there should be enough of a time delay that they are on the verge of forgetting before they next practice. How long that delay should be depends on a number of factors, including the learners and the type of material being studied.

PRACTICE ALONE DOESN'T ENSURE MEMORISATION

The multi-store model emphasises the processes of repetition and maintenance of ideas in working memory as the key route into LTM; the 'residue of thought' idea. If we think about and practice things a lot, it implies, then they should stick in memory.

However, we also know that learners can think about things *a lot* and still fail to remember them. Most teachers have experienced the situation where ideas and skills that were practised intensively as part of previous lessons are apparently absent at the time of an assessment. Likewise, things we could do with ease when we ourselves were at school are not always remembered easily now. Could you get the same grades on your school exams if you sat the papers today? Sadly, it's not uncommon for learners to retain very little of the information and skills covered in a school course, especially if they are tested months or years later.

It is not the amount of study time that matters but how *effective* the methods used during a study session are. Consider, for example, the issue of 'retrieval practice vs re-reading', which has already been discussed. These might take a similar amount of time, but retrieval practice is more effective. Clearly, just holding information in your working memory is not the whole story. It matters what you *do* with that information. Further evidence comes from research on homework by Schumann et al. (1985). They found that hours spent studying homework account for less than 1 per cent of the variance in students' grades.

This means that putting in the hours is a very limited strategy. It's much more important that the hours are being used effectively. Naturally, this is great news for students! They can study for *less* time and get *better* results.

THE ROLE OF ACTIVE LEARNING

A further fundamental limitation of the multi-store model is that it does not account for active learning. Although repetition in working memory is to some extent an active task, we have probably all been in the situation where we repeat a phrase without really focusing our attention on it. For example, you can probably do an everyday task while repeatedly muttering, 'one, two, three; one, two, three ...' under your breath.

Likewise, a learner can listen to a lecture or read something on a whiteboard and screen – and perhaps even copy it into their notes – without really taking it in. It's unlikely that they will remember much of what they wrote down, if anything, when doing this task in a way that requires neither paying full attention to the content, nor actively doing anything with the information.

This is why *active learning* is seen as important. Our definition of this term does not imply that learners need to be physically active. However, they do need to *do* something with the material. Active learning implies that they are using information, thinking about it and processing it in some way, rather than passively listening, re-reading, or copying. So, when planning for learning, it is important to consider what the students will do with the material. Teaching it in small chunks of learning with time to practise built in (as recommended by Rosenshine, 2012) helps to ensure the learning is active.

A DEVELOPED THEORY

We have seen that the multi-store model presents a useful outline, but omits some important details. The *new theory of disuse* is another major theory of LTM and, in our view, a more practically useful one. It will inform the next few sections.

Developed by Robert Bjork and colleagues, the new theory of disuse doesn't look at LTM as a box where things are passively stored and gradually forgotten, but rather sees each memory as having a chance of being retrieved that depends on what we do with it. This is much more applicable to teaching and learning in the classroom. It also places much less emphasis on short-term repetition (quite the opposite, in fact).

STORAGE STRENGTH VS RETRIEVAL STRENGTH

The theory distinguishes between two properties any memory has – storage strength and retrieval strength. Any educational experience or task can affect the chance that a memory is more easily retrieved in future and can affect how securely it is stored. It doesn't necessarily affect both things equally.

Storage and retrieval of memory might sound similar, but we have already looked at how they can be seen as two separate processes in memory (see Chapter 1). In the context of the new theory of disuse, storage and retrieval strength again have a fundamental difference between them:

- retrieval strength relates to how easily we can bring something to mind and therefore affects *performance* in the here and now, or in the near future;
- storage strength relates to how securely something is stored in LTM and therefore affects learning and transfer both now and over the long term.

We can, in fact, put these two things on a diagram, and consider every memory in terms of both storage and retrieval.

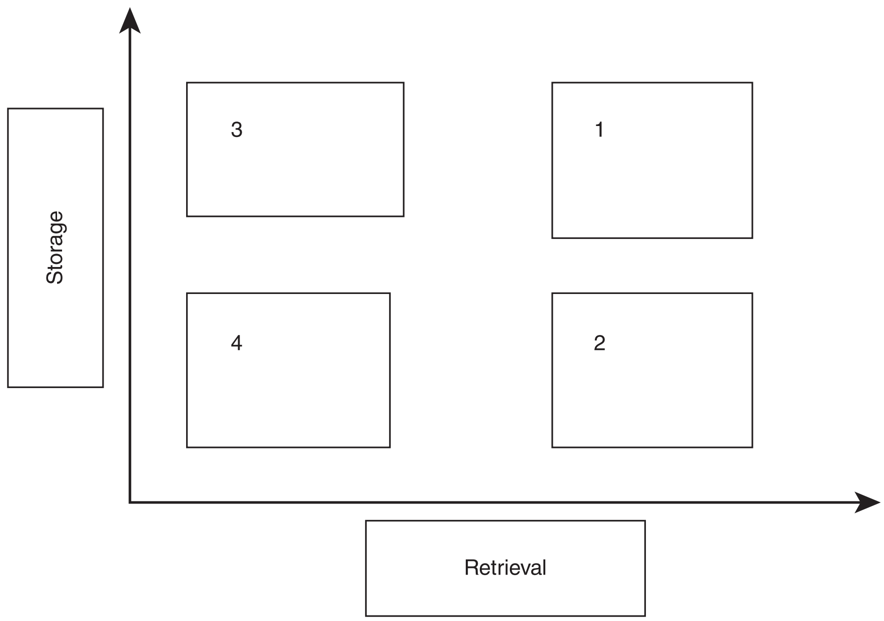

Figure 3.2 Retrieval strength and storage strength of memories

The best-case scenario is clearly that things are *both* easy to retrieve and securely stored (area 1 on the graph). Memories with these features spring easily to mind and can be recalled and used after a considerable delay. But what would it look like if this wasn't the case?

- If a memory has high retrieval strength but low storage strength (area 2 on the diagram), it can be recalled in the near future (hours or days) but will soon be forgotten, like material a learner has crammed the night before a test or exam.
- Typically, a person's memory for their hotel room number has high retrieval strength but low storage strength. Most people can remember their hotel room number during a stay, even after a delay of a few hours while they are doing another task such as sightseeing. It's therefore stored in LTM, not working memory. But can you remember your hotel room numbers from last year, or five years

ago? For most people, the answer is 'no'. This temporary information fades rapidly from memory and is not well stored over the long term.
- If a memory has high storage strength but low retrieval strength (area 3), this means it is in LTM but hard to get out when needed. These are the memories which are really hard to dredge up from your mind. In education, an example might be when a student sits in an exam hall and struggles to remember something. In a more everyday context, something like the postcode of your childhood home or the names of distant relatives might have high storage strength but low retrieval strength. The information is in there, but it probably takes you some time to retrieve.
- If memories have low storage strength and low retrieval strength (area 4), this is the worst-case scenario. These have not been securely learnt, and they don't come to mind easily either. If students didn't engage in effective practice and consolidation during the lesson, then they probably couldn't tell you the key information five minutes later, never mind retrieve it months later in their exam.

THE ROLE OF ORGANISATION

You might wonder why some information which is intensively studied, for example via cramming in the lead up to an exam, has low storage strength. At its worst, cramming leads to taking in a bunch of disconnected facts without understanding how they connect together. The information has high retrieval strength, but only temporarily. Forgetting is already kicking in fast, before the student even gets to the exam hall. How can it be *in* LTM, and then later be gone? And why do some things stick with us, even when we weren't making much effort to learn them?

The answer relates to schemas. These are organising structures in memory (see Chapter 1). If information isn't connected together – meaningfully organised – then the learner ends up with knowledge that isn't sufficiently resilient or flexible. Additionally, the disconnected nature of the facts makes it harder for learners to *think with* the information. A learner can't make comparisons, analogies, or use the information to help them evaluate other ideas, and knowledge is harder to apply (Perkins & Salomon, 1988; Willingham, 2007). These skills (evaluation, application, etc.) are, of course, central to many courses.

We will say much more about the role of meaningful understanding in the next chapter, but it is worth pointing out that for a student to understand information and connect it into schemas makes that material more resilient, and less prone to forgetting. However, it still needs to be consolidated actively, and after a delay.

PERFORMANCE VS LEARNING

Retrieval strength and storage strength are useful, but they remain theoretical ideas and, without more information, it is hard to know what strength a specific memory will have. Soderstrom and Bjork (2015) make the ideas of retrieval strength and storage strength more relevant to the classroom by linking them to two educational ideas we can actually observe:

- *performance*: how well and accurately a learner can do a task in the here and now. Are they fluent and able to avoid errors?
- *learning*: how well and accurately a learner can do a task after a lengthy delay or in a different context.

Good performance, on the one hand, indicates high retrieval strength. It means the learner can do the task in the here and now, and they may well leave the classroom or study session feeling confident. But this can be misleading.

Learning, on the other hand, is underpinned by high storage strength (though high retrieval strength is also very helpful). This is where information is well embedded in memory and can be retrieved and used out of context and after a delay – for example, when learners can use maths concepts years later in a science class.

A key idea the researchers raise in their paper is that performance and learning are negatively correlated. Strategies that improve performance are *worse* for learning:

- better performance in class – worse learning;
- worse performance in class – better learning.

If that sounds bizarre, take a moment to consider some real examples. The counter-intuitive idea makes sense when you reflect back on some of what you have read so far in this book.

Consider, for example, repetition versus retrieval practice. It's easier for learners to copy from a slide or textbook, or to repeat things verbatim. It's harder for them to retrieve these things actively from memory. They will make more errors on a retrieval practice task, but it is better for learning.

The spacing effect is another clear example. A delay will make performance worse, but it is a more effective way of getting things to stick.

These examples show we shouldn't seek or encourage perfect performance in the classroom. A very easy consolidation test would be like lifting a weight at the gym. If it's too easy, it is of little benefit. If pupils are making no mistakes, is it possible their tasks are too easy and don't include enough active, retrieved or delayed elements? At times, practice needs to be more challenging in order to make learning more effective.

This idea is summed up by a term that has become widely used in education over recent years – 'desirable difficulties'. A desirable difficulty is a strategy that makes tasks more difficult for students (leading to more of a struggle and probably increased errors), but in a good way (Bjork & Bjork, 2011) in that it promotes learning rather than performance. Retrieval and spacing are both desirable difficulties and, arguably, most forms of active learning are desirable difficulties too. Further examples will be discussed later in this chapter.

It's tempting as a teacher to think a key part of one's job is to make things as straightforward as possible for the learners during a task. The idea of performance vs learning helps us to reconsider this. We should be aiming to improve learning, not performance.

Another surprising point made by the researchers behind the new theory of disuse is as follows: 'Forgetting can foster learning' (Soderstrom & Bjork, 2015, p. 192). If it seems strange that forgetting would be useful, consider this – that's exactly what is happening with the spacing effect. More time passing means more forgetting is taking place. But this means that future practice is more impactful.

REFLECTION POINT

Take a moment to think about some of the practice and consolidation tasks that you have used. Are they challenging enough? And do you emphasise and encourage perfect performance?

ACTIVITY

Note down three or four ways in which tasks may be made less routine and challenging in your context.

REALISM AND TRANSFER

You might be wondering *why* factors such as retrieval practice and spacing boost learning rather than performance. The researchers behind the new theory of disuse have an elegant explanation for this.

In their 2011 chapter, Robert and Elizabeth Bjork point out that a lot of the strategies which are considered evidence-based mimic real-world learning situations in some way. Consider, for example, how we learn to drive. It wouldn't be helpful to make practice very easy and repetitive, such as driving around the same block again and again. It would lead to good performance, but poor learning, as it wouldn't prepare us for the real thing.

Or think about remembering how to change a tyre or set up a spreadsheet. In the workplace or everyday life, these are not tasks you would do several times in the same day. They are things which happen occasionally, at times unpredictably, and you have to remember what to do from memory. That is to say, you need to retrieve the skills and knowledge actively, in a different context and after a delay.

These examples show that a key consideration in memory-informed teaching is that the practice should, as far as possible, resemble the situation where knowledge and skills are later used. The difficulties which improve learning mimic real-world challenges. When students leave the classroom and use their learning in the real world, they will have to retrieve relevant skills and knowledge, use them actively and do so out of context and after a delay. These same factors make practice more effective. It means the practice was better preparation for the real thing. Thus, it's not just *any* difficulty that we want, but one which prepares the learner for later use and application of what they have learnt.

It's worth focusing in on the idea of the real-world context. Psychologists have known for decades that memories are harder to retrieve in a context which doesn't resemble the learning situation. Even Pavlov's dog responded less strongly to a tone which was higher or lower pitched. For our learners, it's vital they can recognise new contexts and apply what they have learnt. Researchers call this *transfer*.

For example, it's not very helpful if geography students are great at recognising map features in a classroom, but totally unable to recognise them when out hiking or mountain climbing. The skill needs to transfer.

As another example, consider the language skills we hope students will learn in English language classes. The intention is for these to transfer, so the learners become better at using language in their lives more broadly. But does studying a poem in Year 10 really make you better at giving a workplace speech, or writing a technical manual?

Psychologists have come to recognise transfer is actually very difficult. This is well set out by researchers Barnett and Ceci (2002). They explain a system of transfer according to the following principles, each of which vary on a continuum from low to high:

- how similar is the topic/subject (for example, science vs arts)?
- how similar is the physical context (for example, a lab vs outdoors)?
- how similar is the stimulus (for example, transferring facts about a bird's biology to a reptile vs to a worm)?
- how recently did the initial learning take place (for example, yesterday vs last year)?

The researchers discuss these and other factors which make learning easier or harder to transfer. Each one can vary along a continuum; transferring learning in a different laboratory would be harder than in the same laboratory, for example, but easier than outdoors. The factors also combine; a combination of a different topic/subject, location and stimulus, as well as a time delay, would make transfer very challenging indeed!

To some extent we can manage these factors for our students, but, again, it is worth considering *easier isn't necessarily better*. If we practice harder types of transfer (what researchers call *far transfer*), it will make it easier for students in the future. Ways to help student to transfer their learning during the initial practice stages include:

- giving hints and pointing out the similarities to previous tasks – for example, when showing students how to do a complicated mental maths problem, break it down into smaller sums previously that students have learnt how to solve;
- allowing them to try several tasks which have things in common, prompting reflection each time.

Build up the difficulty gradually, moving step by step from a similar task/application to a more distant one and providing scaffolding. For example, in business studies, Mr P would demonstrate the construction of a break-even graph from start to finish to the class so they can see the steps needed to complete the task. This would then be followed up by guided practice where students individually construct a break-even graph in tandem with the teacher demonstrating each step one at a time. Students would then practice collaboratively with reduced teacher input and, finally, reach the stage where they can create a break-even graph independently.

Another way to tackle the difficulty of transfer is to include more variation in student's practice, and this is discussed next.

VARIATION AND COMPLEXITY

A further desirable difficulty is the extent to which practice tasks are *varied*. When practice is varied, learners become more accustomed to recognising how their learning can be applied and used in different ways. Again, this is not something accounted for in the multi-store model of memory, but the new theory of disuse helps us to understand that greater variation makes retrieval harder and storage more secure. The benefits to learning link mainly to overcoming some of the difficulties with transfer discussed above.

Some of the ways teachers can vary a task include:

- changing the academic context – for example, locating a principle within a scenario question, or engaging in project work or cross-curricular learning – for example, applying skills learnt in computing science lessons to another subject when creating digital media;
- using a different skill, such as switching from factual recall to application – for example, in history – instead of trying to recall verbally the facts of the events that occurred in the Hiroshima bombing; students could be asked to create a story board and discuss the role of this event in world peace movements;
- practising in a different physical location, such as learning outdoors, or setting a task as homework – for example, asking students to create a model of a mosque or church at home for their religious education lesson, labelling it to show the different aspects rather than simply drawing it in their exercise book or labelling a diagram;
- changing the social context, such as moving from individual work to group work – for example, students of politics could be asked to write notes on the dangers of fascism, and then discuss this in a group;
- being in a different mood or physical state – for example, in psychology, students might play a humorous revision game towards the end of the week, which would contrast with the more serious work done during practicals.

Delays lead to variation, so, while this is not the only benefit of the spacing effect, returning to an area of the curriculum a few months later will inevitably vary a number of things, such as the mood a learner is in, their level of knowledge and what they have been thinking about lately. Their social situation and the way the classroom looks might also be different.

Variation of practice is especially valuable with the most fundamental aspects of skills and knowledge, as these will crop up time and again. For example, most students will have a need for arithmetic skills, will need to operate a personal computer and will need to write. If such skills are practised in varied ways, they will be easier to apply flexibly when the need arises.

Granted, the future is inherently unpredictable. It can be hard to know *exactly* how our students might need to use a particular skill or fact they learn at school. If anything, though, this means varied practice is probably a good idea for *all* subjects and topics. It won't do any harm to learning, and can make it a lot more resilient if the need arises.

While useful for learning and transfer, variation will tend to negatively impact short-term performance. Just think how much worse you drive if you are in an unfamiliar car, or how much harder it is to play sport or cook with unfamiliar equipment! Again, it's important to make students aware of this. They need to understand that making more errors is not a bad thing – it's part of the process of developing a more resilient skill.

Related to the idea of variation is *complexity*. How complex are the practice tasks we give, in terms of the number of elements or 'moving parts'? A set of short-answer questions is less complex than a scenario question, which may in turn be less complex than an extended project. Of course, it's fine to use practice which is simplified in some way, especially in the early stages. But, at some point, it can help with transfer if tasks become more complex and authentic. If they don't, then they won't fully prepare learners for using their new knowledge and skills in the real world (beyond the exam).

For example:

- children in the early years practice with phonics and simplified texts, but the goal is for them to move on to independent reading;
- students who are new to a sport like basketball will practice movements, throws and footwork in an isolated way, but only so they can eventually put these skills together in the more complex context of a match.

MANAGING DIFFICULTY

As difficulty increases, there has to be some consideration of student emotions, and ideally a bit of ownership on the part the learner – they need to understand and want these challenges. It is often best to make changes one step at a time; sometimes, new learning strategies may be unfamiliar to students. They should be encouraged to persevere, accepting short-term pain for long-term gain.

In the classroom, we often have to make a choice between a harder or a more straightforward practice task. At times, especially when learners are struggling for confidence or are new to the subject/topic, it might make sense to work on simpler tasks and build up to more challenging things (for more on this issue, see Chapter 6, where we discuss the issue of managing learner needs and modifying the difficulty level of various strategies).

Hopefully, if well managed, desirable difficulties will reduce anxiety over time and boost confidence as learners themselves realise they are improving and developing sound knowledge. Agarwal et al. (2014) found evidence that application of desirable difficulties such as retrieval practice reduces anxiety in school pupils, as the more they practice the better they become at retrieval and addressing any misconceptions. As noted by Dweck (2006), encouraging and praising students may lead to a fixation on performance (rather than learning). We should actively push students to take on challenges, and help them to understand that struggling at times is a normal part of learning. This education in 'how to learn' can support your use of desirable difficulties.

REFLECTION POINT

Are there times when your learners seem a bit relaxed and the difficulty could be increased? And are there times when they are overwhelmed and anxious? Could a better understanding of how to learn help them to manage their anxiety?

DIFFERENT STORES OF LTM

There is one further essential aspect of LTM that should be considered before concluding this chapter, and that relates to storage of different *types* of information (e.g., verbal, visual) that we remember. There are different ways to look at these:

- successors to the multi-store model tend to presume LTM has various stores – one for visual information, one for meaningful information and one for life events, among others;
- the idea of schemas suggests memories are interconnected. As conceived by Piaget and others, a schema includes not just factual information, but sensory and movement information (e.g., Piaget, 1950). It is assumed these are connected together into a kind of network.

Both of these ideas probably have at least an element of truth. An item can be stored differently depending on whether it is visual or verbal, and learners can retain something better if they encounter both verbal and visual information. This is called *dual coding* (Clark & Paivio, 1991). A YouTube video showing a historical moment is more memorable than trying to simply describe it or showing a static picture, a worksheet or whiteboard, and this is partly because it uses more than one sensory modality.

By extension, it is going to extend and enrich any form of learning if more sensory information is present in examples. The more examples and different types of sensory information learnt, the better chance the learner has of retaining one of them and thereby retrieving the knowledge that they need successfully. Exclusively using very simple and abstract forms won't be helpful.

At the same time, memories are interconnected. A simple example of this is that if you are shown an image of a piece of fruit, the word for that foodstuff will automatically spring to mind. So even if there are *separate stores* for visual and verbal information, they must be deeply interconnected. Richer knowledge in LTM will also influence processing in WM (see Chapter 2).

On a practical level, this means that we should be aiming to provide rich learning experiences, and that keeping things in very simple and abstract forms won't be helpful. Likewise, while verbal explanations are important; using other formats as part of the learning process is helpful too, such as stories.

Finally, the interconnected nature of memory brings us on to a discussion of the role of meaning. What exactly *is* meaning, and how does our learning become meaningful? What role does meaning play in the way things are retained in LTM – or forgotten? That will be the focus of the next chapter.

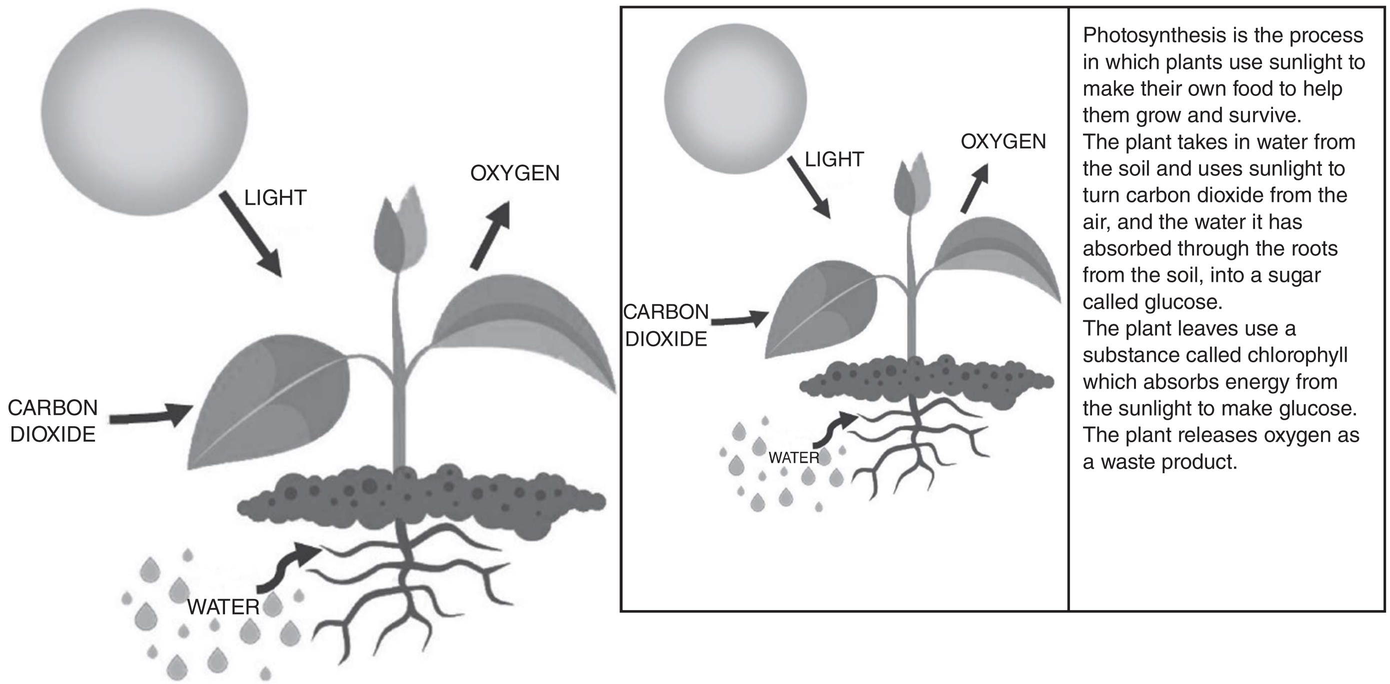

Figure 3.3 An example of dual coding

Note: adapted from Green Plant Sprout image by Vectorportal.com

CONCLUSION

We have looked at some of the key features of long-term memory, focusing on two important theories – the multi-store model and the new theory of disuse. The multi-store model presents a helpful way of showing how LTM connects with working memory, and places an emphasis on how information is taken in and may be lost through forgetting.

However, as we have seen, this model is over simplistic. It assumes that repetition is the main way of getting information into LTM, and treats the two stores as largely separate in their functioning. The multi-store model also doesn't fully account for the spacing effect or active learning – both important factors in getting new learning to stick.

The new theory of disuse helpfully separates out the storage and retrieval strength of a memory, helping us to understand why things may be remembered today or tomorrow in class, but yet rapidly forgotten. That is to say, it emphasises learning rather than performance. To achieve learning, the model supports the use of desirable difficulties such as spacing and retrieval practice. These harm performance and lead to errors, but boost learning over the long term.

Finally, we saw how a key target for educators is that information transfers to other contexts, but that this is far from straightforward. Several strategies for boosting transfer, are discussed, including varying the conditions of practice.

KEY POINTS

- As with working memory, there are disagreements about LTM, including both how it is structured and how new information is encoded, stored and later retrieved or forgotten.
- Repetition is important, but it is not the whole story. Things can be repeated multiple times and still forgotten.
- Retention can be divided into two key components: storage and retrieval. Both are important, and educational strategies that focus on just one (retrieval, for example) are limited. To put it another way, learners need to be able to both retain things over the long term, and retrieve and use them when needed.
- Spacing out practice and making practice active (e.g., via active retrieval and application) helps to make sure information is secure in LTM rather than being rapidly

forgotten. These strategies can be applied to multiple areas of the curriculum (and beyond).

- Simplifying tasks for learner is not always the best way forward. Simpler tasks which are more similar to previous practice and more repetitive will be good for *performance*, but more complex tasks that are more similar to real life and more varied will be good for learning.
- Practice should (as far as possible) resemble the situation and conditions under which knowledge and skills will later be used.