04

Activating hard thinking

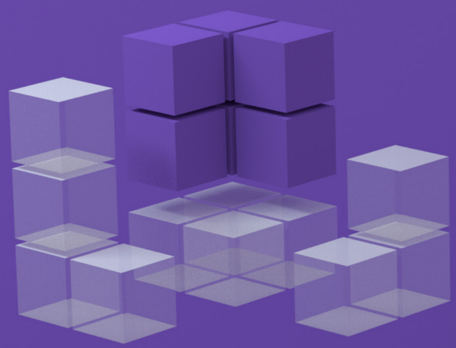

Great teachers present content, activities and interactions that activate their students' thinking

In many ways, Dimension 4 represents the heart of great teaching: getting students to think hard about the material you want them to learn. It may also be the hardest part of the job to learn, partly because it is rare to get reliable feedback about whether it is working: student learning is invisible, slow and non-linear, so how can we tell if it is happening?

Summary of Dimension 4

4.1 Structuring: giving students an appropriate sequence of learning tasks; signalling learning objectives, rationale, overview, key ideas and stages of progress; matching tasks to learners' needs and readiness; scaffolding and supporting to make tasks accessible to all, but gradually removing them so that all students succeed at the required level

4.2 Explaining: presenting and communicating new ideas clearly, with concise, appropriate, engaging explanations; connecting new ideas to what has previously been learnt (and re-activating/checking prior knowledge); using examples (and non-examples) appropriately to help learners understand and build connections; modelling/demonstrating new skills or procedures with appropriate scaffolding and challenge; using worked/part-worked examples

4.3 Questioning: using questions and dialogue to promote elaboration and connected, flexible thinking among learners (e.g., 'Why?', 'Compare', etc.); using questions to elicit student thinking; getting responses from all students; using high-quality assessment to evidence learning; interpreting, communicating and responding to assessment evidence appropriately

4.4 Interacting: responding appropriately to feedback from students about their thinking/knowledge/understanding; giving students actionable feedback to guide their learning

4.5 Embedding: giving students tasks that embed and reinforce learning; requiring them to practise until learning is fluent and secure; ensuring that once-learnt material is reviewed/revisited to prevent forgetting

4.6 Activating: helping students to plan, regulate and monitor their own learning; progressing appropriately from structured to more independent learning as students develop knowledge and expertise

Elements of Dimension 4

Partly because this fourth dimension is so complex, there seems to be a wide range of different ways to present it in different existing frameworks. We have split it into six elements here, though the total weight of content in this dimension means they are each quite broad and inevitably overlapping. It seems likely that when we start to develop instruments to give teachers feedback about their development, some further splitting may be required. Our six elements are: structuring, explaining, questioning, interacting, embedding and activating.

1 **Structuring** refers to the choice, matching and sequencing of learning tasks and signalling how they contribute to learning goals. Great teachers share learning aims with their students in ways that help students to understand what success looks like. This does not mean simply writing out lesson objectives or (worse still) getting students to copy them down. Abstract statements of learning aims may be useful but are certainly not enough. To specify learning aims properly, teachers also need to have examples of the kinds of problems, tasks and questions learners will be able to do, as well as examples of work that demonstrates them, with a clear story about how and why each piece of work meets each aim. Great teachers also help students to understand why a particular activity is taking place and how current learning fits into a wider structure. They draw attention to key ideas and signal transitions between activities that focus on different parts of the journey.

A component of structuring is the selection of learning tasks. Tasks must present an appropriate level of difficulty for each student: hard enough to move them forward, but not so hard that they cannot cope, given the existing knowledge and resources they can draw on. Tasks must also promote deep rather than just surface-level thinking (Hattie, 2012), focusing on abstraction, generalisation and the connectedness and flexibility of ideas rather than just reproduction of facts or procedures. In planning a curriculum, tasks must be sequenced so that prerequisite knowledge and skills are accessible and fluent when they are needed. Great teachers build in opportunities for review to check this is the case – and adapt their plans if not.

Great teachers also recognise that complex tasks often require scaffolding: beginning with a simplified or limited version of the task to make it manageable. This often requires some differentiation, as different learners may begin with different levels of readiness and different capacity for learning new material. A knowledge of individual students' needs, including SEND, comes into play here. However, one of the defining characteristics of great teachers is that they require all students to achieve success (Hattie, 2012). Scaffolding provides a gentler entry, but the destination remains the same. Lower-attainers may take longer and need more help, but the job of teachers is to 'disrupt the bell curve', not just to preserve it (Wiliam, 2018). The crucial thing about scaffolding is that you take it away as ideas and procedures become secure and fluent: by the end, those complex tasks are accessible to all.

2 The second element of Dimension 4 is **explaining**. All teachers present new content and ideas to students, but the best presentations have concise, appropriate, engaging explanations that are just right for the students: neither too short nor too long; neither too complex nor too simple. Evidence from both **cognitive load theory** (CLT, Sweller et al., 1998, 2019) and direct instruction (Adams & Engelmann, 1996; Stockard et al., 2018) supports the importance of good explanations. In presenting material, teachers should pay attention to the 'cognitive load' it presents to their students: limiting the number

Cognitive load theory:

Since Sweller first proposed cognitive load theory in the 1980s, it has evolved. A key thread that runs through this research is that humans' capacity for processing information is limited. Working memory can only handle so much at a given moment; this is dependent on the type and complexity of the information. Understanding CLT can be helpful in becoming better at explaining, but there is more to it.

and complexity of new elements; breaking complex ideas or procedures into smaller steps; helping students to assimilate concepts into – and extend – existing schemas; minimising extraneous, irrelevant or distracting input, from either content or environment. Presentations should be planned, crafted and refined, using the collective expertise of experienced teachers and the wisdom of trial and error, to make them as effective as possible.

Part of the skill of explaining is connecting new ideas to prior knowledge. Great teachers know that durable and flexible knowledge depends on connecting ideas together, creating and modifying schemas. A schema is a cognitive structure that enables information to be organised and stored in long-term memory. Schemas are very powerful for learning because they allow individual bits of knowledge to be 'chunked' together into an overarching principle or concept, or for a series of procedures to be combined into a single 'script', and hence processed as a single element. A simple example would be a beginning reader's schema for the letter 'a', which allows them to recognise that a whole range of different shapes (e.g., a, a, a, **a**, a, A, *A*) are actually equivalent in terms of their meaning. The steps in a procedure, such as column subtraction, or conjugating regular -er verbs in the present tense in French, can also be stored as a schema, allowing the whole process to be treated as a single, automated element that can be drawn on in solving a more complex problem. Prior knowledge is structured in schemas and the process of acquiring new knowledge consists of accommodating it into existing or modified schemas and making connections between them (CESE, 2017; Sweller, 1994). Hence, learning depends on the connections that learners make between new ideas and what they already know. Great teachers activate that prior knowledge, reinforce it and connect new ideas to it.

A key insight here is that long-term memory is not just a storage facility, analogous to an encyclopaedia or information searchable on the internet; nor is it limited to routine facts. Instead, the structure and connections among elements of memorised knowledge are precisely what enable it to be used in solving problems or performing complex tasks: if it is not structured and accessible in memory, it cannot be used. Conversely, if a student has a good store of well-structured knowledge, and fluent, automated skills, absorbing new ideas and procedures is much easier. In the same way that gardeners prepare the soil before sowing seeds, great teachers prepare their students for new knowledge by ensuring their existing schemas are well-connected, fluent and accessible. This is one of the reasons why presenting great explanations is not just a generic skill, like being a good communicator: it depends on a detailed knowledge of the content and ideas being explained and how they are learnt.

One of the ways teachers explain new ideas is with the use of examples. Examples can make the abstract concrete and support conceptual understanding if used appropriately (Booth et al., 2017; Braithwaite & Goldstone, 2015). Examples supply content to the theory-building and schema-developing processes that are necessary for new knowledge to be connected, classified and stored. Also necessary for these processes are non-examples and borderline cases: the exceptions and hard cases that define the boundaries of a rule or definition. For learners to construct strong schemas, they need to understand the limits between what does and does not count as an example.

There is also a good deal of evidence that the use of worked examples can be helpful in introducing new ideas (Booth et al., 2017; Sweller et al., 2019). Particularly effective are 'completion problems' where students are given partial solutions and required to complete them. These can help students to focus on the examples but also manage the difficulty level while retaining authentic tasks.

3 Our third element is **questioning**. Pretty much every model of teaching includes this in some form. For example, Rosenshine enjoins us to 'ask a large number of questions and check the responses of all students' (2010, p. 12). But questioning is already one of the commonest things teachers do, and the key to quality is not the number of questions but the type and how they are used. For Hattie (2012) it is about the balance between deep and surface-level thinking that teachers promote. When Smith et al. (2008) searched for the strongest differentiators between 'expert' and 'experienced' teachers they found a focus on promoting deep learning to be one of five distinguishing characteristics (along with: presenting content effectively; creating a learning climate; monitoring and giving feedback; believing that all students can succeed). Hattie (2012) defines this deeper understanding as 'more integrated, more coherent and at a higher level of abstraction'. The key point is that just asking a lot of questions is not a marker of quality; it's about the types of questions, the time allowed for, and depth of, student thinking they provoke or elicit, and how teachers interact with the responses.

This raises an important distinction between different reasons teachers do questioning. Understanding and promoting great teaching requires us to attend to teachers' purposes as well as their practices: not just what they do, but why they do it; what problems they are trying to solve (Kennedy, 2016). Teachers use questioning for two main – and quite distinct – purposes: to promote students' thinking, and to assess it.

In the former purpose, questioning is a tool to promote deep and connected thinking. Great teachers use questioning as part of a dialogue in which students are engaged and stretched. They prompt students to give explanations and justifications for their answers, or just to improve an initial response, to describe their thinking processes, to elaborate on their answers, exploring implications, 'what-if's and connections with other ideas and knowledge (Dunlosky et al., 2013; Praetorius et al., 2018). Although we have used the word 'questioning' here, the range of activities teachers use to promote oracy and dialogue are much wider. They may also encourage students to ask their own questions. Shimamura (2018) encourages learners to apply the 'three Cs' (categorise, compare and contrast) and 'elaborative-interrogation' (asking, and answering, 'why' and 'how' questions) to help them learn new ideas. Great questioning promotes deep student thinking, helping them to connect and elaborate ideas.

In questioning designed for the latter purpose, the focus is on eliciting and checking student thinking, knowledge and understanding: in other words, assessment. Asking questions, or providing prompts, that provide clear insight into whether students have grasped the required knowledge and understanding is hard; it is in the nature of assessment (and indeed all human communication) that student responses are always equivocal, and interpretations should be probabilistic rather than certain. Questioning that is interactive may go some way to overcome this if follow-ups and prompts are used skilfully to clarify. Great teachers also have strategies for checking the responses of all students. Asking meaningful and appropriate questions that target essential learning, collecting and interpreting a response from every student, and responding to the results, all in real time in the flow of a lesson, is hard to do well, but great teachers do it and it is probably a skill that can be learnt.

Whether questions are asked interactively or as part of a fixed assessment process, starting with great questions that provide maximum information is key. When used for the purpose of assessment, questions should be seen as tools to elicit insights into students' thinking. Questions provide information if they discriminate between those who know and those who don't yet. Whether an assessment is a single question or a formal examination, great teachers understand the amount of information it provides, how much weight it carries and what inferences and decisions it can support. They understand that what has been learnt is not the same as what has been taught (Nuthall, 2007) and that assessment is the only tool we have to make the former visible, albeit 'through a glass, darkly'. Crucially, they plan and adapt their teaching to respond to what assessment tells them.

4 This responsiveness is at the heart of our fourth element, **interacting**. The quality of learning interactions between teachers and students is central to the learning process. Interactions may be seen as a form of feedback, and again there are two distinct purposes here: feedback to teachers that informs their decisions, and feedback to students that helps them learn.

The former purpose, feedback to inform teacher decisions, overlaps considerably with the previous element. Information from questioning and assessment is the basis of this feedback. But it is how the teacher responds to the feedback that matters. First of all, teachers have to understand and interpret the assessment result appropriately. They may need to check or verify that their interpretations are correct. They also need to appraise the context accurately, being sensitive to the needs, history and dispositions of the student(s) involved. Then they need to identify and decide among a set of options for action. Each will have trade-offs between, for example, time, effort and reward. If some students need more time and help with a topic while others are ready to move on, for example, this may be a hard choice. Finally, they need to implement the chosen option effectively to achieve the desired learning.

For the latter purpose, feedback goes the other way: to the student. Although we know that feedback can enhance learning powerfully (Hattie & Timperley, 2007), we also know that the mediating effects of different combinations of kinds of feedback, learner and task characteristics and different ways of giving feedback are extremely complex. There is no simple recipe for giving powerful feedback. Feedback can help by clarifying or emphasising goals or success criteria ('Where am I going?', Hattie & Timperley, 2007), thus directing students' attention to productive goals. It may draw attention to a gap between actual and desired levels of performance ('How am I going?'), which, again, may be positive if goals are challenging, accepted and accompanied by feelings of self-efficacy (Locke & Latham, 2002). It may cue attributions for success or failure to reasons the student can control, such as effort or strategy choice (Dweck, 2000). Or it may indicate productive next steps ('Where to next?', Hattie & Timperley, 2007). This last mechanism may be the hardest to predict and deliver, precisely because it is a complex interaction between what the learner knows already, what they need to know and their readiness to do what is required to bridge the gap. It also requires an expert judgement about the kinds of actionable next steps that are most likely to deliver the most learning, given all these variables. Great teachers have enough knowledge and experience of similar situations to develop sound intuition about what is likely to work best (Hogarth, 2001), but such intuition is hard to capture in simple rules.

5 The fifth element is **embedding**, getting the learning to stick. The importance of embedding learning rests on the insight from cognitive load theory that memory is not just a storage facility for facts that could just as easily be looked up: the schemas that we use to organise knowledge in memory are the very things we use to think with and to connect new learning to (Sweller, 1994).

There are numerous ways great teachers embed learning. One is by ensuring that students practise any procedures that are regularly required to be fluent and accurate. A large body of psychological research shows that 'overlearning' (continuing to practise after performance has reached a specified standard) can be important for producing learning that is durable and flexible (Soderstrom & Bjork, 2015). Knowledge or schemas that are required for future learning must be secure and readily retrievable. Forgetting is normal but can be slowed or prevented by periodic revisiting and review.

Great teachers ensure that students practise until learning is fluent, automatic and secure.

An important point to note here is that student practice generally needs to be monitored and guided initially (Rosenshine, 2010). In new learning, there is typically a transition: practice begins as helping to learn the ideas, developing connections and understanding, and building schemas; then follows consolidation, gaining confidence and fluency, in which scaffolds and other supports are removed, as is the need for teacher guidance and monitoring; finally comes embedding, where practice becomes independent, fluent, accurate and automatic. Great teachers understand and plan for this transition, monitoring and supporting each student's passage through it and ensuring there is adequate time for each stage.

Practice is particularly effective if it is distributed or 'spaced' over time, with deliberate gaps between for forgetting. Distributing practice like this makes learning feel harder and reduces performance during actual practice, even though it is more effective in the long term – what Bjork and Bjork (2011) have called a 'desirable difficulty'. Great teachers provide opportunities for students to practise procedures and recall of information that must be learnt until it is fluent, and to repeatedly revisit and re-practise after allowing time to forget.

Another approach to embedding is to exploit the 'testing effect', requiring learners to generate answers or recall information from memory in a (low-stakes) test-like process. Again, a vast body of research shows that this is the single most effective way to increase long-term retrieval strength: the ability to recall information or procedures after a delay (Adesope et al., 2017). Moreover, the benefits of testing are not limited to simple recall; the process of having to search for and generate answers also strengthens the connections with, and retrievability of, related information (Delaney et al., 2010). As with all learning, students get better at what they are required to do, so it is important to require them to answer questions that go beyond simple recall and surface-level thinking. Great teachers use the testing effect to delay forgetting with questions that require deep and connected thinking. And of course, testing and spacing can be combined by making time to revisit previously learnt, but about to be forgotten, material after a suitable delay.

There are also other practices that, if done well, can help to ensure learning is durably and flexibly embedded. These include interleaving, varying the conditions of practice, elaboration, and self-explanation (Bjork & Bjork, 2011; Dunlosky et al., 2013; Weinstein et al., 2018). Great teachers understand the principles behind these effects and the contexts in which they are likely to be useful, have a range of strategies for deploying them in practice, and incorporate appropriate and effective use into their teaching.

Metacognition:

Although it has a simple literal meaning of "thinking about thinking," metacognition has developed into a broad umbrella term for a number of related cognitive processes. Different frameworks have chosen to focus on different aspects or definitions of this concept. Ultimately, the associated strategies share the aim of helping learners plan, monitor, and evaluate their learning.

6

The sixth and final element of Dimension 4 is **activating**: helping students to become independent by planning, regulating and monitoring their own learning. Activating, and in particular promoting, student **metacognition**, is a feature of many of the research-based frameworks (e.g., Ko et al., 2013; Praetorius et al., 2018; van de Grift et al., 2017).

When teachers introduce new ideas, it is appropriate to be directive: presenting structured content explicitly, directly teaching what needs to be understood. However, for most educators, the larger aim is to wean students off this dependency on the teacher, encouraging them to become independent, self-actualised learners. In some accounts, this contrast is presented as a polarised opposition between 'traditional', teacher-led, didactic approaches on the one hand, and, on the other, 'progressive', student-focused, constructivist methods and beliefs. In part at least, this division reflects a misunderstanding of the complexity of teaching: different approaches work best at different times, with different students, according to different learning aims, at different stages in the learning process, etc. One approach doesn't fit all.

Within cognitive load theory, both the 'expertise-reversal effect' and the 'guidance-fading effect' refer to the finding that strategies such as presenting limited, structured content and worked examples, which work best for 'novices' (i.e., students who do not yet have the knowledge of the topic or domain encoded in schemas in long-term memory) are no longer the most effective for 'experts', whose chunking and automation of individual elements allow them to tackle and learn more from solving whole problems (Sweller et al., 2019). Using problem-solving as a teaching strategy is overwhelming and inefficient for learners who do not have the required background knowledge, but becomes optimal and necessary when they do.

Interventions to promote the use of metacognitive strategies are among those with the largest effects on attainment, and strategies to help students plan, monitor and evaluate should be explicitly taught and supported (EEF, 2018). Students of all ages should be explicitly taught strategies to plan, monitor and evaluate their learning, ideally in the context of the specific content they are learning. Great teachers also draw attention to their own planning and self-regulation when they model the process of completing complex tasks, and similarly encourage students to 'self-explain' their thinking.

Evidence for Dimension 4

Dimension 4 of our model is derived from the 'cognitive activation' dimension of Praetorius et al. (2018), but it features in every other framework too.

A significant challenge with this dimension is that most of the teacher behaviours that have been found to be effective for activating students' thinking are quite complex. There isn't a simple recipe for developing students' metacognition, for example, or for giving students actionable and appropriate feedback. The very same teacher action could be good in one context, with students of a particular age, history and level of knowledge/skills, in relation to a particular piece of work, and bad in another. Some of the instruments for evaluating teaching make a distinction between 'low-inference' indicators, where the judgement or rating can be made quite easily and 'high-inference', where a more complex specification and a considerable amount of training for raters is required to get consistency. For many of the really powerful elements of cognitive activation, capturing a valid indicator is at the 'high-inference' end of this continuum. It may also be that even expert, trained observers simply cannot perceive enough of the complexity and subtlety of the classroom context to make valid judgements about whether a particular practice is 'good'. This certainly feels like an area where verbal descriptions of practice are inadequate, or at least only a starting point. Understanding what each element means and what really excellent practice looks like could be seen as a life's work.

Another complexity is prioritising among all these elements. Not all of these are important for every teacher to work to improve. It may be, for example, that some parts of the previous three dimensions are prerequisites for this one: if you don't have the content knowledge, or basic classroom management, then those should come first. Some elements of cognitive activation may be a career-long project: even an experienced, expert teacher may find value in improving these aspects of their practice. Some may be best bets for quite a large group of teachers. Wiliam (2018) argues, for example, that a small number of strategies within this dimension, grouped as comprising formative assessment, are likely to offer the highest leverage for most teachers. Rosenshine's (2010) ten principles of instruction may also be seen as high-leverage skills within this dimension.

We think the jury is still out on this question of priorities: existing evidence and theory cannot give an individual teacher a clear-enough steer about which element they should prioritise, or even whether they should try to work on more than one. In the subsequent stages of this project we hope to collect data from teachers working in different ways to improve their practice so that we can learn how to match different kinds of advice, guidance and support to the individual needs of a teacher in order to have the biggest positive impact on student learning.

Structuring is an explicit focus of many of the existing frameworks. For example, it is one of the eight dimensions of the Dynamic Model (Creemers and Kyriakides, 2011), as is 'orientation', which involves clarifying and sharing objectives, and is merged here under the heading of structuring. Careful curriculum sequencing is emphasised in a number of well-validated models of teaching, including mastery learning and direct instruction (Creemers et al., 2013). Selection of appropriate learning tasks and matching their difficulty to students' existing knowledge and readiness, including scaffolding for difficult tasks, features in many models and reviews of effective instruction (e.g., Ko et al., 2013; Muijs et al., 2018; Praetorius et al., 2018; Rosenshine & Stevens, 1986; van de Grift et al., 2017). The need for a balance between foundational knowledge and higher-level extension into 'deep thinking' for all learners is also widely supported (e.g., Hattie, 2012; Pianta et al., 2012; Praetorius et al., 2018).

Support for the importance of explaining draws on evidence from both cognitive load theory (Sweller et al., 1998, 2019) and direct instruction (Adams & Engelmann, 1996; Stockard et al., 2018), as does the use of examples, non-examples, worked examples and completion problems (Booth et al., 2017; Braithwaite & Goldstone, 2015; Sweller et al., 2019). The importance of clear presentation of ideas is an explicit focus of both the ISTOF and ICALT frameworks (Muijs et al., 2018; van de Grift et al., 2017).

Teachers' use of questioning is also widely featured in the evidence-based frameworks (e.g., Creemers & Kyriakides, 2011; Muijs et al., 2018; Rosenshine & Stevens, 1986; van de Grift et al., 2017). Most of these emphasise the importance of the types of questions asked and how teachers respond to them, as do Hattie (2012), Smith et al. (2008) and Ko et al. (2013). The use of elaborative interrogation is judged to have 'moderate utility' by Dunlosky et al. (2013). Questioning as part of formative assessment has a strong evidence base (e.g., Wiliam, 2010).

Interacting denotes the quality of learning interactions between teachers and students, including feedback in both directions. Evidence for the importance of feedback in learning is abundant (e.g., Hattie & Timperley, 2007; Kluger & DeNisi, 1996), especially if combined with goal-setting (Locke and Latham, 2002). Evidence about the role of feedback in prompting adaptive attributions has been cited above under Dimension 2, Element 4.

Embedding learning through practice and retrieval features in some frameworks (e.g., Creemers & Kyriakides, 2011; Rosenshine, 2010) but is noticeably absent from others. These practices draw both theoretical and empirical support from cognitive science, including studies in authentic school classrooms (e.g., Adesope et al, 2017; Delaney et al., 2010; Dunlosky et al., 2013; Weinstein et al., 2018).

Activating – supporting students as self-activated learners – includes strategies that encourage independence, planning, regulating and monitoring. These teacher behaviours are explicitly mentioned in many of the research-based frameworks (e.g., Ko et al., 2013; Praetorius et al., 2018; van de Grift et al., 2017). Extensive evidence from intervention studies supports explicit teaching of metacognitive strategies (e.g., Donker et al., 2014; Hacker et al., 2009).