the **next** word|

The Mental Parliament

Distributed cognition, large language models, and the limits of knowing your student

PAPER Thought Piece

PREPARED BY Matt Esterman, Founder & CEO, The Next Word, in partnership with Claude.ai, Google Gemini, and Google's NotebookLM.

DATE April 2026

WEB nextwordnow.com

A truly interesting audio conversation created in NotebookLM is available at nextwordnow.com

A premise we have not examined

For decades, education has worked from a quiet assumption: that a student is a coherent, knowable individual. We build profiles. We map characteristics. We design differentiated tasks. We write reports that describe a child as a consistent entity moving through our classrooms, as if the sum of observed moments adds up to a single, stable person.

The model we have been working with does not survive contact with the science.

The brain is not a single voice deliberating. The brain is, in David Eagleman's phrase, a team of rivals: many networks, competing for the same output channel, producing the behaviour we then attribute to a unified self. Eagleman puts the case directly:

> *As Walt Whitman correctly surmised, we are large and we harbor multitudes within us. And those multitudes are locked in chronic battle. There is an ongoing conversation among the different factions in your brain, each competing to control the single output channel of your behavior. As a result, you can accomplish the strange feats of arguing with yourself, cursing at yourself, and cajoling yourself to do something, feats that modern computers simply do not do.*
>
> David Eagleman, Incognito: The Secret Lives of the Brain

This paper argues that the implications of that picture reach beyond neuroscience and are about to be reshaped again by artificial intelligence. Recent peer-reviewed work shows that large language models share computational principles with the brain in ways that go beyond metaphor. Both are distributed probabilistic systems. Both produce output by sampling from competing internal tendencies, weighted by context. The implications cross four positions in any school community: the teacher, the leader, the parent, and the student themselves.

The question this opens is more uncomfortable than the AI integrity debate the sector is currently having. Can we genuinely know a student, in the way the Australian Professional Standards for Teachers requires, if that student is a parliament rather than a person?

What the brain actually is

The neuroscience here is not fringe. It is the consensus of the last forty years of cognitive science.

Michael Gazzaniga, the neuroscientist whose split-brain research transformed our understanding of consciousness, has described the architecture of the brain in terms that make any unitary model of selfhood difficult to sustain:

> *The brain has millions of local processors making important decisions. It is a highly specialized system with critical networks distributed throughout the 1,300 grams of tissue. There is no one boss in the brain. You are certainly not the boss of the brain. Have you ever succeeded in telling your brain to shut up already and go to sleep?*
>
> Michael Gazzaniga, Who's in Charge? Free Will and the Science of the Brain

The contemporary neuroscience of large-scale brain networks has moved this from metaphor to mechanism. Vinod Menon's triple network model, one of the most-cited frameworks in current systems neuroscience, describes cognition as the dynamic switching between three competing networks: a default mode network running self-referential and internally directed thought, a central executive network handling externally focused attention and working memory, and a salience network gating between them. Cognition functionally is the negotiation among these systems, and disorders ranging from depression to schizophrenia to ADHD are increasingly understood as failures of network coordination rather than localised malfunctions.

Stanislas Dehaene, in his global workspace theory of consciousness, gives the picture its sharpest expression. Most of what the brain does happens in parallel, unconsciously, across many specialised processors. Consciousness is the moment when one of those processes wins access to a shared broadcast channel:

> *When we say that we are aware of a certain piece of information, what we mean is just this: the information has entered into a specific storage area that makes it available to the rest of the brain. Among the millions of mental representations that constantly crisscross our brains in an unconscious manner, one is selected because of its relevance to our present goals. Consciousness makes it globally available to all our other processes.*
>
> Stanislas Dehaene, Consciousness and the Brain

The conscious self that a teacher addresses, the part of a student that responds to instruction and reflects on feedback, is one process among many. It is the broadcast, not the system. The other processes continue running, with their own goals, their own assessments of threat and reward, their own contributions to behaviour. They are also the student.

A note on terminology. The team of rivals framing belongs to Eagleman as a popular shorthand. The technical literature speaks of distributed processing, large-scale brain networks, and predictive systems. The framing matters more than the label. A human being is a parallel system whose output emerges from competition and coordination among many parts.

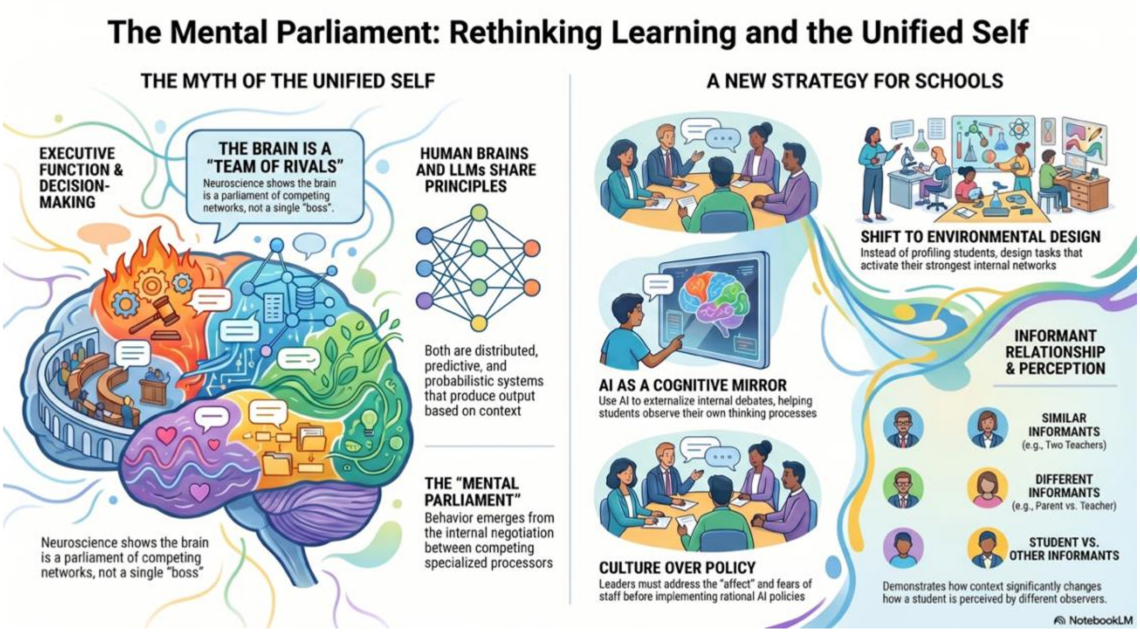

The artificial mind in the same family

This is where the AI conversation in schools needs to grow up.

When we talk about ChatGPT or Claude in a classroom, we tend to talk about them as tools. Useful or threatening, accurate or hallucinatory, helpful or corrosive to learning. What we rarely acknowledge is that the architecture producing those outputs has begun to look uncannily like the architecture producing our own.

The peer-reviewed evidence here has moved fast. In 2021, a team led by Martin Schrimpf at MIT, working with Evelina Fedorenko and Nancy Kanwisher, published one of the most consequential findings in recent cognitive neuroscience:

> *The most powerful models predict neural and behavioral responses across different datasets up to noise levels. Models that perform better at predicting the next word in a sequence also better predict brain measurements, providing computationally explicit evidence that predictive processing fundamentally shapes the language comprehension mechanisms in the brain.*
>
> Martin Schrimpf et al., Proceedings of the National Academy of Sciences, 2021

Up to noise levels is the crucial phrase. Modern transformer-based language models predict human neural activity during language processing about as well as the brain's own variability allows. The work has been replicated and extended, including in a Nature Neuroscience paper by Ariel Goldstein and colleagues at Princeton, using direct intracranial recordings from human cortex:

> *We provide empirical evidence that the human brain and autoregressive deep language models share three fundamental computational principles as they process the same natural narrative: both are engaged in continuous next-word prediction before word onset; both match their pre-onset predictions to the incoming word to calculate post-onset surprise; and both rely on contextual embeddings to represent words in natural contexts.*
>
> Ariel Goldstein et al., Nature Neuroscience, 2022

Andreas Doerig and a group of leading computational neuroscientists, in a 2023 Nature Reviews Neuroscience paper, give the field a name. They call it neuroconnectionism, and they describe artificial neural networks as occupying:

> *A Goldilocks zone of computational abstraction, which allows them to model complex cognitive functions grounded in sensory data while still being mappable to biological features.*
>
> Andreas Doerig et al., Nature Reviews Neuroscience, 2023

The honest qualifier matters. The same Princeton group that demonstrated the convergence is careful to note its limits:

> *Unlike humans, deep language models cannot think, understand or generate new meaningful ideas by integrating prior knowledge. They simply echo the statistics of their input. Going beyond the importance of language as having a central organizing role in our cognition, deep language models indicate that linguistic competence may be insufficient to capture thinking.*
>
> Goldstein et al., Nature Neuroscience, 2022

A defensible position, then, is partial convergence in language processing rather than architectural identity. LLMs and human brains are not the same thing. They share a family of computational principles. They are both distributed, predictive, probabilistic systems whose outputs emerge from competition and weighted combination across very many parameters. When you prompt a model, you are not querying a database. You are activating a particular region of a probability distribution, and the same model will produce different outputs depending on what you make salient.

This has direct relevance to the classroom. We teach students to use AI as if it were a search engine with a personality. Helping them understand it as a distributed system whose output depends on what you activate would change how they prompt, what they expect, and how they reason about the answers they receive. It would also give them a working model for understanding what is happening inside their own heads.

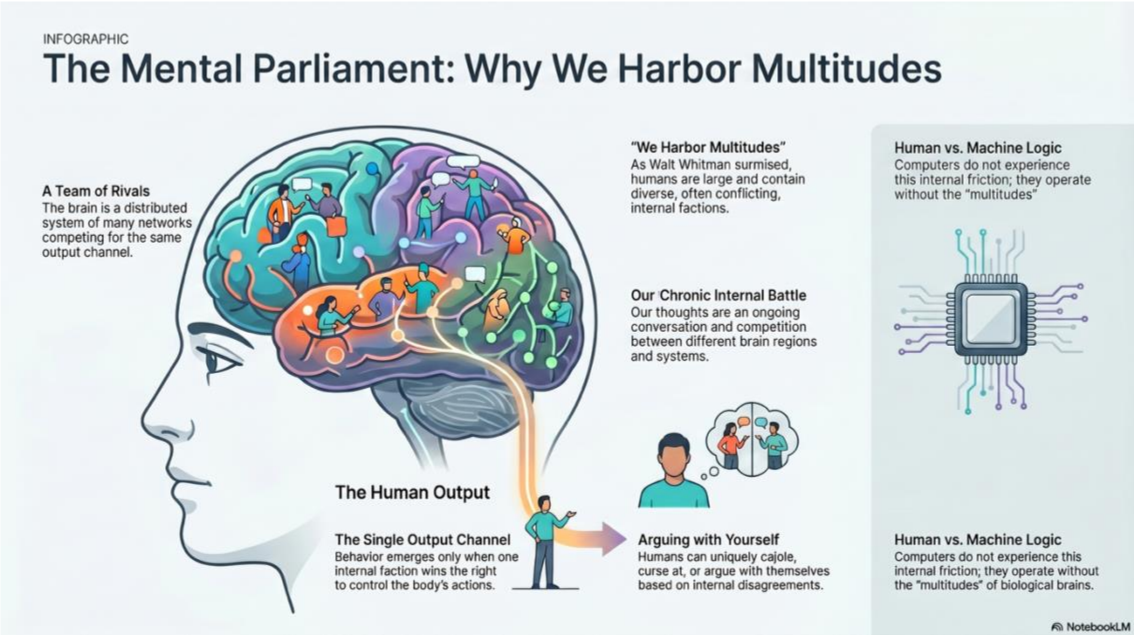

The teacher's position

Standard 1 of the Australian Professional Standards for Teachers requires teachers to know students and how they learn. At graduate level, this means demonstrating knowledge of physical, social and intellectual development, of how students learn, and of strategies for differentiating to meet a full range of needs.

This is a reasonable aspiration. It rests on an epistemological claim that becomes harder to sustain the more closely we examine it: that a student has a learnable, stable profile that can be known and applied across contexts.

Mary Helen Immordino-Yang, drawing on a decade of affective and social neuroscience research, has put the situation as plainly as anyone in education has:

> *Teachers intuitively know that neither their nor their students' learning is steady and constant, the same day in and day out and moment to moment, consistent from topic to topic. Rather, we all have good and bad days; moments of excitement, engagement and inspiration and moments of*

> *disappointment, disengagement and frustration. These differences influence how children learn and how teachers teach; they even affect what students know at a given time. In short, learning is dynamic, social and context-dependent because emotions are, and emotions form a critical piece of how, what, when and why people think, remember and learn.*
>
> Mary Helen Immordino-Yang, Emotions, Learning, and the Brain

Immordino-Yang and Antonio Damasio's earlier formulation goes further. The aspects of cognition that schools care most about are not optional add-ons to a rational core. They are subsumed within the affective system itself:

> *The neurobiological evidence suggests that the aspects of cognition that we recruit most heavily in schools, namely learning, attention, memory, decision making, and social functioning, are both profoundly affected by and subsumed within the processes of emotion; we call these aspects emotional thought.*
>
> Mary Helen Immordino-Yang and Antonio Damasio, Mind, Brain, and Education, 2007

Reinhard Pekrun's control-value theory of achievement emotions, the most empirically grounded account of how affect shapes learning, makes the same point in different language. The emotions a student experiences in a learning task are not stable traits. They are appraisals, generated in response to perceived control and perceived value within a specific context. Change the context, change the appraisal, change the emotion, change what the student can access cognitively. Brown, Collins and Duguid argued from a different angle as far back as 1989 that knowing itself is inextricably situated in the physical and social context of its acquisition and use, and cannot be extracted from these without being irretrievably transformed.

What follows for the teacher is uncomfortable but liberating. The student who is confident in small-group discussion is running a different network configuration from that student in a whole-class cold-call situation. The student who reads carefully at home is not operating the same cognitive architecture when their phone is in reach. The child who is creative in art and anxious in mathematics may have a single distributed system that activates differently under different conditions.

Teachers know this. The phrase they were having a bad day is a recognition that a different part of the student showed up. Our professional frameworks have not given teachers language for this. We still write the comment about the whole student, as if a stable description were possible.

A more honest practice would shift the question from what does this student need to what conditions activate this student's strongest networks. It would mean designing learning environments rather than student profiles. It would mean treating engagement as a system response rather than a student attribute.

The implication for differentiation is uncomfortable in a different way. If students are distributed systems, and those systems respond differently to context, affect, social environment, and perceived threat, then differentiation at the level of individual neural networks is not tractable at scale. The cognitive architecture works against it. John Sweller and his colleagues, three decades into cognitive load theory, have established that working memory is the binding constraint. It is true for the student trying to learn, and it is true for the teacher trying to read thirty students simultaneously. The empirical record on differentiation as commonly practised is more mixed than the rhetoric suggests. Robert Slavin's best-evidence synthesis on within-class ability grouping reported a median effect size of zero. Holly Hertberg-Davis, writing from inside Carol Tomlinson's own intellectual lineage, offered an unusually candid assessment:

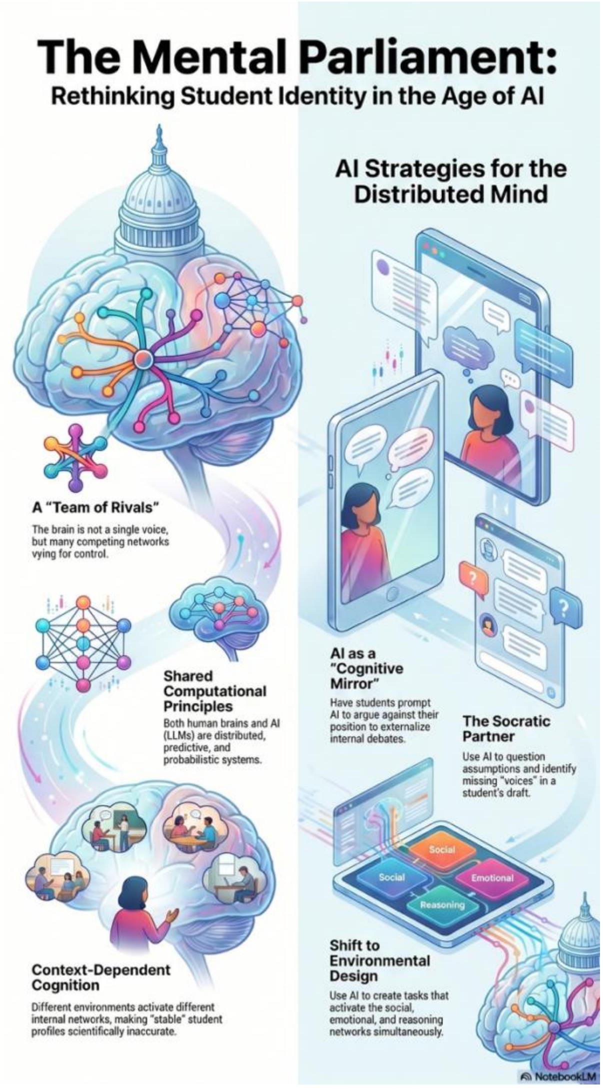

It does not seem that we are yet at a place where differentiation within the regular classroom is a particularly effective method of challenging our most able learners.

Holly Hertberg-Davis, Gifted Child Quarterly, 2009

What is functionally implausible at the level of individual cognitive trajectories may still be possible at the level of conditions. A teacher cannot consciously track thirty internal parliaments through fifty minutes. A teacher can design a task whose stakes, audience, social structure and emotional texture activate more of each student's competing networks at once. Real audiences. Genuine choices. Things that matter socially and emotionally. Tasks where the habit-brain, the social-belonging-brain, the threat-avoidance-brain, and the reasoning-brain all have something in play. That is a more complete learning environment, and it is the most honest answer to the differentiation question that the cognitive architecture allows.

The leader's position

School leaders run a version of this problem at scale, and in two directions at once.

The first is internal. The leader who presents a confident, unified front is not unified inside. They are managing a continuous internal negotiation among the voice that wants to protect their people, the voice that knows the system must change, the voice that fears getting it wrong in public, and the voice that genuinely believes in the school's potential. The best coaching has always worked at this level, even when it did not use this language. Helping a leader notice which internal voice is currently driving is more useful than analysing the decision they made.

The second is institutional. A staff body is itself a parliament, and a school's culture is the emergent result of dozens of internal negotiations interacting under particular environmental conditions. Andy Hargreaves named the structural omission in change theory more than two decades ago, and the omission persists:

> *One of the most neglected dimensions of educational change is the emotional one. Educational and organizational change are often treated as rational, cognitive processes in pursuit of rational, cognitive ends. If emotions are acknowledged at all, this is usually in a minimalist way in terms of human relations or climate setting, where the task of leadership is to manipulate the mood and motivation of their staffs, in order to manage them more effectively. The more unpredictable passionate aspects of learning, teaching and leading, however, are usually left out of the change picture.*
>
> Andy Hargreaves, International Handbook of Educational Change, 1998

The popular synthesis of this insight, drawn from Jonathan Haidt's rider-and-elephant metaphor, is that change efforts that address only the rational brain produce understanding without motivation. Chip and Dan Heath put it this way:

> *If you reach the Riders of your team but not the Elephants, team members will have understanding without motivation. If you reach their Elephants but not their Riders, they'll have passion without direction. In both cases, the flaws can be paralyzing.*
>
> Chip Heath and Dan Heath, Switch

AI governance programs in schools often fail at exactly this point. A new policy is presented as a rational necessity. The reasoning brain hears it and agrees. The habit-brain still grades the same way. The threat-brain is still nervous about being caught not knowing something. The social-belonging-brain is still watching what colleagues are actually doing. Policy without culture work addresses one voice in the staff parliament while leaving the others untouched. Recent RAND research on US districts implementing AI training found that the more effective districts began with teachers' fear and discomfort rather than with instructional tools. Affect first, policy second.

The leaders doing the most credible work on AI in Australian schools are not the ones with the most polished policies. They are the ones holding the multiple voices in their staff body without forcing premature alignment. Resistance is not an obstacle to manage.

It is information about which networks have been activated, and what conditions would need to change for a different internal conversation to occur.

The parent's position

Parents experience the most disorienting version of the problem, because they are routinely told their child is someone they have never met.

The empirical evidence here is among the cleanest in developmental psychology. Thomas Achenbach's 1987 meta-analysis remains the foundational reference, and its central finding has been replicated for nearly four decades:

> *We found 269 samples in 119 studies for meta-analyses of Pearson rs between ratings by parents, teachers, mental health workers, observers, peers, and the subjects themselves. The mean rs were .60 between similar informants, .28 between different types of informants, and .22 between subjects and other informants. The modest correlations between informants indicate that child and adolescent problems are not effectively captured by present-versus-absent judgments of problems. Instead, the variations between reports by different informants argue for assessment in terms of multiple axes designed to reflect the perceived variations in child and adolescent functioning.*
>
> Thomas Achenbach et al., Psychological Bulletin, 1987

The headline number is r = 0.28 between parents and teachers describing the same child. Andres De Los Reyes and colleagues replicated the finding across 341 newer studies in 2015. The disagreement is robust, and it is not noise. It is signal. The child shows different versions of themselves in different contexts, because different environments activate different parts of the same distributed system.

Urie Bronfenbrenner, the founder of ecological systems theory, captured the methodological consequence with characteristic bluntness:

> *Much of developmental psychology has consisted of the science of the strange behavior of children in strange situations with strange adults for the briefest possible periods of time. To understand the way children actually develop, we must observe their behavior in natural settings, while they are interacting with familiar adults over prolonged periods of time.*
>
> Urie Bronfenbrenner, The Ecology of Human Development

Susan Harter's longitudinal work on adolescent self-development showed that the appearance of multiple selves across contexts is not a developmental problem. It is a normative cognitive achievement of mid-adolescence, and the contradictions between role-related selves are themselves a focus of healthy adolescent reflection.

For parents, this changes the meaning of the disagreement with the school. The child who is withdrawn at school and confident at home is not concealing their true self in one location. They are genuinely both. The home environment activates different parts of the same parliament. The school environment activates others. The information from each setting is real. The mistake is treating either as the whole.

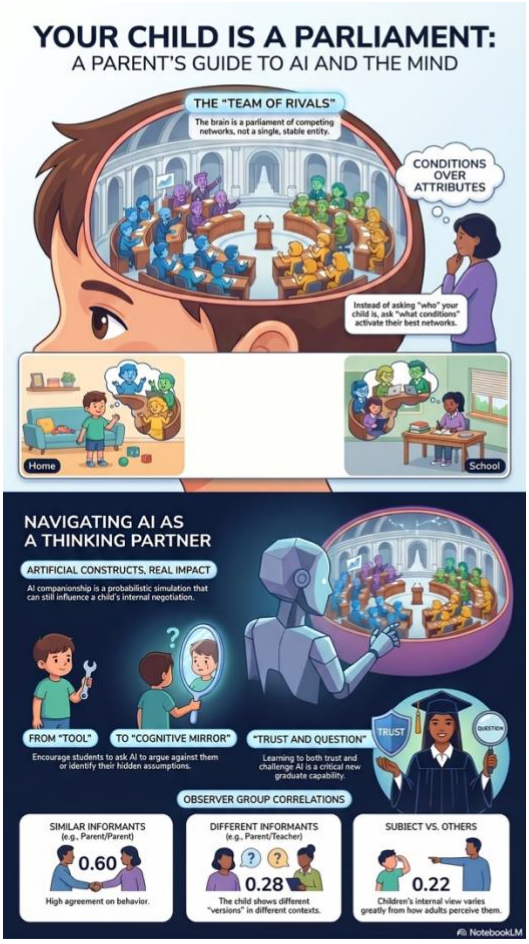

For schools, the implication is structural. A teacher who sees a child for fifty minutes a day, in a single subject, across a class of thirty, is sampling a small region of that child's distributed system. The professional humility to say I have observed this student in these particular conditions is more accurate than the impulse to summarise. Robert Pianta has made the case for years that the relational nature of teacher knowledge is the core of teaching itself, and that the system has consistently underinvested in it:

> *Teaching children involves more than knowing your subject and being able to maintain control of a class. Every child that comes into the classroom is unique. All have different home environments, ethnic backgrounds, abilities, and learning styles. And every child develops a different relationship with the teacher. Yet all too often teachers who are experts in their subjects and experts in general principles of instruction fail to reach students because they do not know how to connect to them.*
>
> Robert Pianta, Enhancing Relationships Between Children and Teachers

Parents are also distributed systems. The parent who presents as calm and engaged in a meeting may be managing an internal negotiation between the voice that wants to trust the school and the voice that needs to protect their child. Dismissing the composed face as agreement is a common and costly mistake. The parent who appears disengaged is rarely uninterested. They may have had their own difficult school experiences, activating networks of threat or shame that make school environments feel adversarial rather than collaborative. Better information rarely solves this. Different conditions sometimes do.

AI is opening a new front in the parent conversation. Within any parent community, you have networks activated by safety concerns, networks activated by competitive anxiety about their child's readiness, networks activated by genuine curiosity, and networks activated by distrust of institutions. A parent communication strategy that speaks to one of these voices is missing most of the room.

The student's position

The most important and least examined question in all of this is what it means for a student to know themselves.

Stephen Fleming, whose laboratory at University College London has produced the most rigorous recent work on the cognitive science of self-awareness, frames the stakes in a way that connects directly to the AI conversation:

> *Many of our legal frameworks, those that ascribe liability and blame for errors, are based on the notion of being able to justify and defend what we did and why we did it. Without an explanation, we are left with blind trust, in each other, or in our machines. Ironically, some of the highest performing machine learning algorithms are often the least explainable. In contrast, humans are rapacious explainers of what we are doing and why, a capacity that depends on our ability to reflect on, think about, and know things about ourselves, including how we remember, perceive, decide, think, and feel. Psychologists have a special name for this kind of self-awareness: metacognition.*
>
> Stephen Fleming, Know Thyself: The New Science of Self-Awareness

Metacognition is well established as a strong predictor of learning outcomes. Most school-based metacognition work asks students to reflect on their thinking as if their thinking were a single stream to observe. The deeper move, and the one most aligned with what the science actually shows, is helping students recognise that they are running multiple competing processes simultaneously, and that much of what they experience as procrastination, distraction, or emotional difficulty is a poorly resolved internal negotiation.

A student who cannot start an assignment is rarely lazy. A network activated by the threat of failure may be winning its competition against a network that wants to engage, because the conditions of the task have made threat-avoidance stronger than approach. A student who understands this can respond to their own behaviour with environmental redesign rather than self-criticism. Where am I working? What is around me? Which voice is loudest right now, and what would change which voice gets the floor?

This is where AI becomes genuinely interesting, rather than threatening or auxiliary.

Andy Clark and David Chalmers' extended mind thesis, first published in 1998, argued that when a human organism is coupled with an external resource in a two-way interaction, the system as a whole is the cognitive system, not the brain alone. In 2025, Clark himself published a paper extending the framework to current generative AI:

> *We humans are and always have been what New York University philosopher David Chalmers and I call extended minds, hybrid thinking systems. Instead of acting as mind-extending technologies, the fear is that these may act as mind-replacing technologies. How might we react to that kind of worry? Here, the devil will (for some time at least) remain in the details. But there is suggestive evidence that what we are mostly seeing are alterations to the human-involving creative process rather than simple replacements. Learning how to both trust and question our best AI-based resources is one of the most important skills that our evolving educational systems will now need to install.*
>
> Andy Clark, Nature Communications, 2025

The student who uses an AI to get a faster answer is using a powerful tool in the least interesting way. The student who uses an AI to externalise their own internal parliament is doing something genuinely new. Ask the model to argue against your position. Ask it to represent the perspective of someone who disagrees. Ask it to identify the assumption hiding inside your question. Ask it to tell you which voice in your draft is loudest and which is missing. The model's distributed nature, its ability to surface different positions on demand, becomes a scaffold for the student's own distributed cognition. The metacognitive practice that good education has always tried to teach finally has an externalisation tool that matches the structure of the mind it is trying to develop.

This is consistent with the most current Australian thinking on AI and cognition. Jason Lodge and Leslie Loble, in their 2026 report for the Australian Network for Quality Digital Education, set out a series of pedagogical patterns that move directly in this direction, including AI as a cognitive mirror in which students teach the model and reveal their own understanding, and AI as a Socratic partner that questions rather than answers. Loble's framing of the educational imperative is the right one for Australian schools now:

> *The educational imperative is not to protect students from a world where AI is the norm, but to prepare them for it. Young people need strong knowledge foundations and complex reasoning skills, not just the ability to Google it or outsource the thinking to AI. By supporting the teacher, we empower the human expert who is best placed to manage the complex, relational work of co-regulating learning, managing cognitive load, and building the evaluative judgement, self-regulated learning, and metacognition that students need.*
>
> Leslie Loble, in Lodge and Loble, Artificial Intelligence, Cognitive Offloading and Implications for Education, 2026

The combination of mature metacognitive practice and intelligent use of AI as a thinking partner is, for this generation, the most important graduate capability schools can build. It is also the natural extension of what the cognitive science has been telling us about the mind for forty years.

Can we know a student?

The AITSL Professional Standards require teachers to know students and how they learn. Standard 1 sits at the foundation of the framework. The graduate-level descriptors ask teachers to demonstrate knowledge of physical, social and intellectual development, of how students learn, and of strategies for differentiation.

The neuroscience of distributed cognition does not make this requirement irrelevant. It makes it more precise, and more honest.

If a student is a parliament rather than a person, then what teachers can actually know is not a fixed profile but a set of patterns. Which environments activate this student's threat-avoidance networks. Which contexts surface the exploratory network. What social configurations produce engagement rather than shutdown. What is the relationship, for this student, between emotional state and access to reasoning.

These are knowable. They require observation across conditions rather than the formation of a stable summary profile. They require curiosity rather than categorisation. They require teachers to hold multiple, sometimes contradictory, observations about the same student without resolving them prematurely into a single account.

They also require honesty about limits. A teacher who sees a student for fifty minutes a day, across a class of thirty, in a single subject, is not in a position to know that student in any deep sense. They are sampling a small region of a complex system. The Grattan Institute's 2022 survey of approximately 5,000 Australian teachers found that more than 90% report not having enough time to prepare effectively, and many report inadequate support for the relational work of helping struggling students. The conditions for knowing students at the depth Standard 1 implies are not present in most Australian classrooms. The standard, taken seriously, is a structural critique of the system in which it is being applied.

John Hattie's well-known synthesis identifies teacher-student relationships as among the strongest influences on student learning. The methodology of Visible Learning has been critiqued, and the effect-size literature should be cited with care, but the relational anchor is robust across more rigorous syntheses. Robert Pianta's longitudinal work on early teacher-child relationships and Debora Roorda's meta-analyses on engagement consistently identify relational quality as a primary mechanism for learning outcomes. Knowing students is not a peripheral activity that supports instruction. It is the substrate on which instruction depends.

What this suggests for AITSL, and for the teacher professional learning ecosystem more broadly, is that the standard of knowing students needs a richer underlying model of what a student is. A model that accounts for context dependence, internal competition, and the irreducibly partial nature of any professional observation. The standard would not need to change in its language. It would need to change in its assumed object.

What this means in practice

For teachers, the shift is from student profiling to environmental design. Ask which conditions activate each student's strongest networks, and build those conditions into the classroom rather than chasing a stable description of the child. Treat engagement as a system response rather than a student attribute. When a task fails, redesign the environment before redescribing the student.

For leaders, the shift is from policy to culture. Take the internal negotiations of your staff body seriously as change material. Surface the competing voices rather than suppressing them. AI governance documents matter, and culture work matters more. Address the affect of your staff before you address the architecture of your policy.

For parents, the shift is from contradiction to context. The child who is one person at home and another at school is reporting accurate information about how their distributed system responds to different environments. The job is not to reconcile the versions but to understand what each reveals about the conditions the child is navigating. Both schools and parents are sampling small regions of a complex person, and a partnership across those regions is more useful than a competition between them.

For students, the shift is from self-criticism to self-observation. Treating self-knowledge as the work of observing a system rather than identifying a fixed identity changes the kinds of questions a young person can ask about their own behaviour. Why does this voice get loud under stress. What conditions help the working voice win the floor. How can I redesign my environment to support the version of me I want to bring to this task. The student who can do this work, with or without an AI thinking partner, is developing the capability that matters most in the world they are entering.

A final word on what we are building

The arrival of AI in classrooms has been framed primarily as a capability and integrity challenge. Can students use it. Should they. What does it mean for assessment.

These are real questions. They are not the deepest ones this moment opens.

The deeper question is whether AI can help teachers, leaders, parents, and students engage more honestly with the complexity of human minds, including their own. Whether prompting an AI to surface multiple competing responses can model and support the kind of internal plurality that good thinking actually requires. Whether AI governance in schools can be designed as an invitation to a richer model of learning and personhood, rather than only as risk management.

The schools doing the most interesting work with AI right now are not the ones with the best policies. They are the ones asking what kind of thinking they want to cultivate, and using AI as a lens on that question rather than as a threat to be managed.

The standards, the policies, and the professional frameworks will catch up. The teachers and leaders who are already working at this level, who hold their students with curiosity rather than certainty, who surface competing voices in their teams rather than demanding alignment, and who are honest about the limits of what any individual can know about any other individual, are already ahead of where the sector is pointing.

The team inside the room is more complex than the room has acknowledged. The parliament inside each student is in continuous session. The work of education has never been about knowing one student. It has always been about creating the conditions in which the better voices in each parliament are more likely to rise, more often, in the company of others doing the same.

That is the work AI can support, and it is the work the science has been describing all along.

Sources

Achenbach, T. M., McConaughy, S. H., & Howell, C. T. (1987). Child/adolescent behavioral and emotional problems: Implications of cross-informant correlations for situational specificity. Psychological Bulletin, 101(2), 213–232.

Australian Institute for Teaching and School Leadership (2018). Australian Professional Standards for Teachers. Education Services Australia.

Brown, J. S., Collins, A., & Duguid, P. (1989). Situated Cognition and the Culture of Learning. Educational Researcher, 18(1), 32–42.

Bronfenbrenner, U. (1979). The Ecology of Human Development: Experiments by Nature and Design. Harvard University Press.

Clark, A. (2025). Extending Minds with Generative AI. Nature Communications, 16(1), Article 5990.

Clark, A., & Chalmers, D. (1998). The Extended Mind. Analysis, 58(1), 7–19.

De Los Reyes, A., Augenstein, T. M., Wang, M., et al. (2015). The validity of the multi-informant approach to assessing child and adolescent mental health. Psychological Bulletin, 141(4), 858–900.

Dehaene, S. (2014). Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts. Viking/Penguin.

Doerig, A., Sommers, R. P., Seeliger, K., Richards, B., et al. (2023). The neuroconnectionist research programme. Nature Reviews Neuroscience, 24, 431–450.

Eagleman, D. (2011). Incognito: The Secret Lives of the Brain. Pantheon.

Fleming, S. M. (2021). Know Thyself: The New Science of Self-Awareness. Basic Books.

Gazzaniga, M. S. (2011). Who's in Charge? Free Will and the Science of the Brain. Ecco.

Goldstein, A., Zada, Z., Buchnik, E., et al. (2022). Shared computational principles for language processing in humans and deep language models. Nature Neuroscience, 25(3), 369–380.

Hargreaves, A. (1998). The Emotions of Teaching and Educational Change. In A. Hargreaves et al. (Eds.), International Handbook of Educational Change. Kluwer.

Harter, S., Bresnick, S., Bouchey, H. A., & Whitesell, N. R. (1997). The development of multiple role-related selves during adolescence. Development and Psychopathology, 9(4), 835–853.

Hattie, J. A. C. (2009). Visible Learning: A synthesis of over 800 meta-analyses relating to achievement. Routledge.

Heath, C., & Heath, D. (2010). Switch: How to Change Things When Change Is Hard. Broadway Books.

Hertberg-Davis, H. (2009). Differentiation in the Regular Classroom Is Equivalent to Gifted Programs and Is Sufficient. Gifted Child Quarterly, 53(4), 251–253.

Hunter, J., Sonnemann, J., & Joiner, R. (2022). Making time for great teaching. Grattan Institute.

Immordino-Yang, M. H. (2015). Emotions, Learning, and the Brain. W. W. Norton.

Immordino-Yang, M. H., & Damasio, A. (2007). We Feel, Therefore We Learn: The Relevance of Affective and Social Neuroscience to Education. Mind, Brain, and Education, 1(1), 3–10.

Lodge, J. M., & Loble, L. (2026). Artificial Intelligence, Cognitive Offloading and Implications for Education. UTS / Australian Network for Quality Digital Education.

Menon, V. (2011). Large-scale brain networks and psychopathology: A unifying triple network model. Trends in Cognitive Sciences, 15(10), 483–506.

Pekrun, R. (2006). The Control-Value Theory of Achievement Emotions. Educational Psychology Review, 18(4), 315–341.

Pianta, R. C. (1999). Enhancing Relationships Between Children and Teachers. American Psychological Association.

Schrimpf, M., Blank, I., Tuckute, G., Kauf, C., Hosseini, E., Kanwisher, N., Tenenbaum, J., & Fedorenko, E. (2021). The neural architecture of language: Integrative modeling converges on predictive processing. Proceedings of the National Academy of Sciences, 118(45), e2105646118.

Slavin, R. E. (1987). Ability grouping and student achievement in elementary schools: A best-evidence synthesis. Review of Educational Research, 57(3), 293–336.

Sweller, J. (2011). Cognitive Load Theory. In J. P. Mestre & B. H. Ross (Eds.), Psychology of Learning and Motivation, Vol. 55. Elsevier.

Tuckute, G., Kanwisher, N., & Fedorenko, E. (2024). Language in Brains, Minds, and Machines. Annual Review of Neuroscience, 47, 277–301.

• • •

About the author. *Matt Esterman is the Founder and Director of The Next Word, an AI strategy and governance consultancy working with K-12 schools across Australia. He works with independent, Catholic, and government sector schools on the governance, strategy, and responsible implementation of AI in education.*

About this paper. *This paper draws on neuroscience, cognitive science, educational research, and recent computational neuroscience work on the convergence between large language models and human language processing, applying that synthesis to the practical work of schools. This has been a genuine, fascinating collaboration between the author and AI tools including Claude, Google Gemini and Google's NotebookLM.*

© 2026 The Next Word · nextwordnow.com

This paper may be shared with attribution.