



Co-inference and collaboration: empowering higher education staff and students through a GenAI literacy development framework based on mind-metaphors

Manish Malik

Canterbury Christ Church University, U.K.

Abstract

Higher education staff and students are beginning to emerge from the liminal space created by the exponential rise in generative artificial intelligence technologies such as large language models. Both groups exhibit a low or developing understanding of the issues related with the use of such tools. On one hand, staff in many institutions are concerned about critical skills development, academic integrity, ethical implications and the long-term cognitive effects of the use of such tools. Equally, students are concerned about the acceptability of the use of such tools at university and the need to develop work-related GenAI skills. On the other hand, foundational large language models are being trained on larger and better data sets and have evolved into 'reasoning' models that can allegedly 'think'. Furthermore, with the emergence of agentic AI, systems that can 'act' independently, claims of shifting agency away from humans to machines are common. This has implications for the world of work and for higher education institutions. This opinion piece presents a new framework that I have developed for enhancing the GenAI literacy of staff and students and to empower and augment themselves with GenAI tools in a balanced, meaningful and collaborative way, whilst maintaining their agency. The mind-metaphors framework includes techniques such as mind-surfing, mind-mending, mind-bending, mind-gaming, mind-storming and mind-stretching, which can help staff and students to develop trust in their chosen GenAI models and master co-inference and collaborative GenAI use.

Keywords: GenAI literacy; co-inference; mind-metaphors; human agency in LLM use; augmentation; GenAI in education.

Introduction

Developments in the field of generative artificial intelligence (GenAI), in particular large language models (LLMs) such as ChatGPT, Gemini, Claude or Perplexity, and more recently, large 'reasoning' models (LRMs) and agentic AI systems, have created a liminal space (Nikolic et al., 2024; Singh and Ngai, 2024; Andersen et al., 2025). Emerging empowered from this liminal space is critical in order to avoid inappropriate use and/or overuse of GenAI tools. Cognitive offloading and an increased reliance on GenAI tools can lead to cognitive debt, diminished learning skills, impoverished long-term memory, weakened problem solving skills and reduced critical skills (Kosmyna et al., 2025; Lee et al., 2025; Li, Liu and Dong, 2025; Oakley et al. 2025). This article introduces a novel mind-metaphors framework designed to empower GenAI users as they emerge from the liminal space created by these technologies. It paves the way for enhancing GenAI literacy alongside developing and exerting user agency through meaningful co-inference and collaboration with GenAI models.

Working in collaboration with a 'GenAI-mind' (non-conscious, with no emotional intelligence or desires) involves aligning your own experience and domain expertise with the GenAI tool's knowledge to meet your intended goals. This has the potential to accelerate productivity while ensuring that users remain in control and retain their agency. This can help actively mitigate the dangers mentioned above, such as cognitive offloading and diminished critical skills.

Beyond the size and quality of the training data, the usefulness of a response in earlier LLMs (e.g. ChatGPT3.5, ChatGPT4.5 or even ChatGPT5 'Instant' mode) is a function of the user's GenAI literacy, their prompting skills and their prior knowledge, which informs their prompt creation. Before diving into mind-metaphors, it is important to understand how LLMs work and how the human/machine agency balance has evolved over time (see Table 1). In a prolonged conversation with an LLM, co-inference refers to the conscious (by humans) and algorithmic and probabilistic (by LLMs) selection and pruning of what stays in its context window affecting future content generation. A context window is defined as a finite number of tokens, in the chat history of a current conversation, that are used in the generative text production process as the chat progresses. Tokens are part of the LLM vocabulary and are full or part words (see next section for examples) which the LLM responds to and generates its responses as. Tokens, represented as numbers, are

mathematically processed by LLMs to generate more tokens. At each step, many possible output tokens with different probabilities exist to choose from. These probabilities are based on the training data used to train the model. An LLM chooses a token to output to the user based on its temperature settings. It may choose something far less probable and be more creative if its temperature is set to high values. At the lowest temperature settings, an LLM chooses the most probable tokens to favour accuracy, although inaccuracies may still be present depending on the training data, the prompt used, and what is already in the context window. The mind-metaphors framework taps into these probabilities to be more or less creative, as will be explained later.

Prompts and their outputs fill the context window as a conversation builds. This is then used by the model to determine the next output token when prompted again. Thus, exerting your agency to control what remains in the context window, helps align what is generated with your actual needs. In other words, past prompts and responses that are permitted to stay in the context window by the user, are used by the models to predict future tokens, forming the sentences and paragraphs in response to a user prompt. Specific mind-metaphors, as will be shown later, empowers you to regain your agency.

Simply put, foundational LLMs, such as ChatGPT 3.5 or the more recent 4.5 or ChatGPT5 Instant, are quick next token predictors. Humans recognise things quickly, through associations and resemblances. Likewise, LLMs use the tokens in their context window to rapidly generate the next tokens, chosen from a list of probable options, forming sentences and paragraphs iteratively. Such automation is analogous to system 1 or fast thinking, as popularised by Nobel laureate and economist, Daniel Kahneman (2011). System 1 thinking requires little or no effort. It is intuitive, fast and automatic in humans, and the underlying mathematics and computations make it work in a similar fashion in LLMs. Foundational LLMs allow the user the most agency (see Table 1), but the user must know how and when to exert their agency when using these models. This is where the mind-metaphors come in handy, as will be shown later.

Kahneman also describes a more deliberate, effortful and therefore slow human process called system 2 thinking. Rapid evolution in the GenAI space, has seen models process prompts more in a system 2 'thinking' way (Hua and Zhang, 2022). LLMs or 'reasoning' models make choices on behalf of the user, rely less on user's prompting skills or knowledge, and more on test-time compute, also known as inference-time compute (Snell

et al., 2024). Test-time compute or inference-time compute refers to the computational resources used by LRMs to autonomously produce their final outputs as they ‘think’ through multiple chains of token generation to bring you its best output. The LRMs or ‘reasoning’ models take time and compute in generating intermediate outputs that fill the context window, and these are then used cumulatively to produce their final outputs. This gives the illusion that these models can deliberate and debate internally and are ‘thinking’ models that work like system 2 thinking. This does, however, take agency away from users as the ‘thinking’ and deliberations are not user controlled, as LRMs determine the outputs less collaboratively. There are contested findings that question the usefulness of LRMs and their likening to human thinking (Opus and Lawsen, 2025; Shojaee et al., 2025). The mind-metaphors give the user a way to con-infer and be in control of the ‘thinking’ they want the model to do for them, which is otherwise missing when using ‘thinking’ models.

Table 1 shows changes to human and machine agency with recent developments in the models, using the example of evolution in OpenAI’s models (Achiam et al., 2023; Cascella et al., 2024), updated to be correct at the time of publication. It shows the general trend using the specific example of OpenAI’s models, though some models may aim to preserve human agency as well. It also lists an agentic AI tool, namely Deep Research from OpenAI, which can make use of tools such as web search to complete its allocated task. Deep Research autonomously searches the web for relevant content and collates a long report, with minimal or no human input post the initial prompt. It attempts to preserve some human agency by offering some choice it deems acceptable to the user. The agentic tool ‘reads’ and attempts to ‘understand’ the relationship between the many concepts present in the search results related to the user request from multiple online sources and its own training.

Table 1. OpenAI models: impact on agency over time

Date	Technology	Human Agency	Machine Agency	Model Type
Nov-22	OpenAI ChatGPT 3.5	High: Humans direct flow of conversation and interpret	Low: Human-like conversations	Foundational
Mar-23	OpenAI ChatGPT 4	High: Humans direct flow of conversation and interpret	Moderate: Multimodal capabilities expand AI's autonomy	Foundational
Nov-23	OpenAI releases	Moderate: Humans define tasks and do result	High: Extended context allows for more complex	Foundational

	GPT-4 Turbo	validation	tasks	
May-24	OpenAI releases GPT-4o	Moderate: Humans guide flow and application and interpret and validate complex outputs	Very High: Advanced capabilities across text, voice and vision	Foundational
Sep 24	OpenAI ChatGPT o1	Low: Humans formulate problems and interpret result	Very high: Advanced problem-solving and logical reasoning	Reasoning
Jan-25	OpenAI ChatGPT o3-mini	Moderate: Humans guide application and validate complex outputs	Very high: Sophisticated reasoning in a more compact form	Reasoning
Feb-25	OpenAI Deep Research	Moderate: Humans provide initial prompts and verify results	Very high: multi-step research and synthesis	Agentic
Feb-25	OpenAI ChatGPT 4.5	High: Humans direct flow and interpretation	Very high: Largest model yet with improved pattern recognition and connections	Foundational
Aug 2025	ChatGPT 5	Low to high: Based on prompt complexity and version, High in 'Instant' mode and Low in Auto, 'Thinking' and 'Pro' modes.	Low to high: Based on prompt complexity and version, Low in 'Instant' mode and High in Auto, 'Thinking' and 'Pro' modes.	LLM, LRM or Agentic AI are chosen dynamically or by user.

The world of work is increasingly interested in the efficiency savings offered by such agentic AI technologies for automating many knowledge and work processes (Dubey, Astvansh and Kopalle, 2024; Faruqi et al., 2025). For example, many edtech startups are increasingly interested in automating educational tasks fully with agentic AI systems (Concannon et al., 2023). As these armies of agents or silicon slaves, in the shape of agentic AI systems, enter the workflows of higher education staff, they can take their agency away from them. This may leave them feeling frustrated and disenfranchised, particularly when presented with the need to evaluate agentic work, which may not meet their exact needs. This misalignment could result in a waste of time and money for higher education institutions. Using mind-metaphors, staff can augment themselves with GenAI and collaborate on completing work that is aligned with their own needs, values and domain of expertise much faster.

Research shows that the best outcomes in using GenAI comes from domain experts, who also know how to use GenAI tools to augment their ways of working (Senoner et al., 2024). Regardless of the advances in the GenAI space, the dependence on what is in the context window for the generation of final outputs remains true for LLMs, LRMs and agentic AI systems. Highlighting the need to control the contents of the context window, can help rebalance this shift of agency from humans to machines, thereby empowering users. For staff, knowing how to retain their agency involves developing and exercising their critical skills and developing an understanding of, and trust in, these models through use. It is through such co-inferring that a more meaningful collaboration can take shape. Equally important is creating a learning and teaching environment that empowers students, so that they can develop, not just their domain expertise, but also their agency in using GenAI meaningfully and critically.

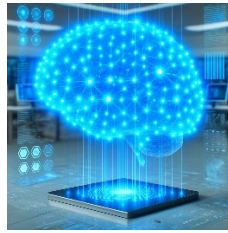
We have indeed come a long way in a short span of time with these models, and yet we continue to see more developments. Whilst GenAI use by staff and students has gone up, GenAI literacy training for staff and students still needs to catch up (Freeman, 2025). The GenAI mind-metaphor framework can help both staff and students to emerge empowered out of the liminal space. Metaphors help link the abstract and unfamiliar concepts through concrete, already understood and familiar concepts (Lakoff and Johnson, 2008). The guiding metaphor useful here is the GenAI as a mind metaphor. To avoid misinterpretation of this, I would like to stress that the 'GenAI-mind' is neither conscious nor has emotions; it is a processing machine, much like a calculator, but for concepts and ideas. The popular models are designed to agree with the user to gain their trust, or in other words show signs of sycophancy (Abels et al., 2025). With critical use of GenAI models and the principles linked to the mind-metaphors framework, staff can generate content that aligns better with their needs and with fewer, or easier to spot, inaccuracies. Using selected mind-metaphors and being aware of the need to be critical, students can invoke different types of learning and control the outputs of the GenAI to help them learn topics of interest or brush up on their metacognitive skills. The next section looks at each of the mind-metaphors and their use for staff and students.

Mind-metaphors framework for co-inference and collaboration with GenAI tools: an explanation

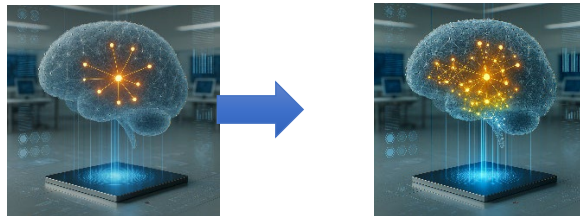
Let us assume that we have generated outputs from a GenAI model, related to all the key concepts that it was trained on, through a prompting exercise. This may not be technically or practically possible, but let us run with this idea for now, to hypothetically populate the context window with all the knowledge a GenAI model has. This now represents a vast set of inter-connected networks of concepts, or schemata, to use a Piagetian term. Figure 1a is the representation of such a 'GenAI-mind' with all its interconnections, like how we may connect and organise ideas in our minds. On a more practical note, the LLM schema related to a concept held in the model, can be externalised to a user and made available into the context window through an appropriate starter prompt, see Figure 1b for a visual representation of this concept. Humans can relate to this knowledge elicitation process as they are familiar with how they and other humans converse and 'pick' each other's brains to find areas of knowledge to collaborate on.

The metaphor of the 'GenAI-mind' gives users an idea of how to use a starter prompt (see below) to first access relevant parts of the schemata which were built during training and then use it in different ways to achieve their objectives. Prompts inspired by the mind-metaphors, give the users greater agency to co-infer and collaborate by managing and processing the contents of the context window to meet their objectives. The LRMs on the other hand, using Test-time compute or inference-time compute, also attempt to take more time and 'think', but they do so in an automated and non-collaborative fashion (Snell et al., 2024). There is little or no opportunity for a user to co-infer with the tool. The mind-metaphors seek to address this loss of agency, and in the process, empower the user in using LLMs with agency.

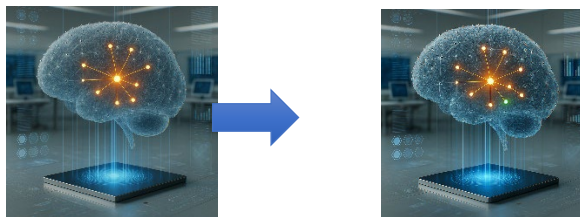
Figure 1. The GenAI-mind-metaphors visualised (generated by ChatGPT4o).



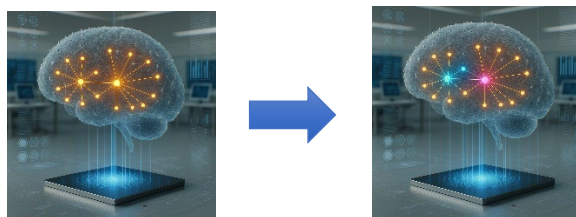
(a) The 'GenAI-mind' shown with its entire schemata.



(b) Externalising a schema before mind-surfing and exploring topics using mind-surfing.



(c) Controlling the context window through mind-mending prompt.



(d) Generating in expected and unexpected ways using mind-bending prompt.

Rebalancing the agency between the user and the tool helps avoid the problems of misplaced expectations, overreliance and reduced critical thinking related to anthropomorphising (Akbulut et al., 2024; Shardlow and Przybyła, 2024). One should not lose sight that it is a user’s agency that matters, and the metaphor of the ‘GenAI-mind’ used here alongside other mind-metaphors are there as guides for relating with and effective GenAI use. Let me explain the role each of the mind-metaphors play in empowering the user in detail next.

Mind-surfing

A foundational LLM, trained on a large natural language corpus or dataset, can produce responses or communication in natural language. Cognition and communication are linked

closely (Sperber and Wilson, 1987; Coulardeau, 2018; and Roberts and Roberts, 2020). Whilst it is true that cognition births communication and vice-versa in humans, such communication can reveal the knowledge structures within the 'GenAI-mind'. Mind-surfing is the act of externalising a specific knowledge schema (the schema in the 'GenAI-mind' relevant to a user) and exploring it using the exact same tokens that the model outputs to enhance the fidelity of its outputs. By externalising and using the same tokens that were outputted by the model, you control the input to the model, such that the underlying mathematics supports the prediction of more relevant outputs. Mind-surfing is, therefore, a process of externalising and more reliably exploring a smaller part of the LLM's schemata (see Figure 1b) that are related to what you want to learn or work with. Mind-surfing is best done within a foundational LLM. However, it can be used before switching to an LRM or agentic AI system, to populate the context window with high-fidelity tokens to control the LRM or agentic AI system outputs to increase reliability.

In LLMs, the output communication depends on the next token prediction capabilities of the Generative Pretrained Transformer (GPT) technology through a mechanism called 'attention' (Vaswani et al., 2017). The attention mechanism is the very mechanism that forces the LLMs to identify the interrelations between the tokens in the context window at any point in a conversation. After determining the important tokens, future tokens are generated iteratively. If a user asks a complex prompt in one go, they are effectively handing over all the agency to the LLM, which uses whatever there is in the context window and the attention mechanism to produce a response it deems suitable for the complex prompt. However, by using a starter prompt (see below) on a chosen broad topic, forces the LLM to use the attention mechanism and output its knowledge structure, informing the user on the most probable output containing its schema for a chosen broad topic 'X'. This empowers the user with agency in what and how to more reliably explore within a topic X that a model has been trained on. The following is one such prompt:

What is there to learn about X?

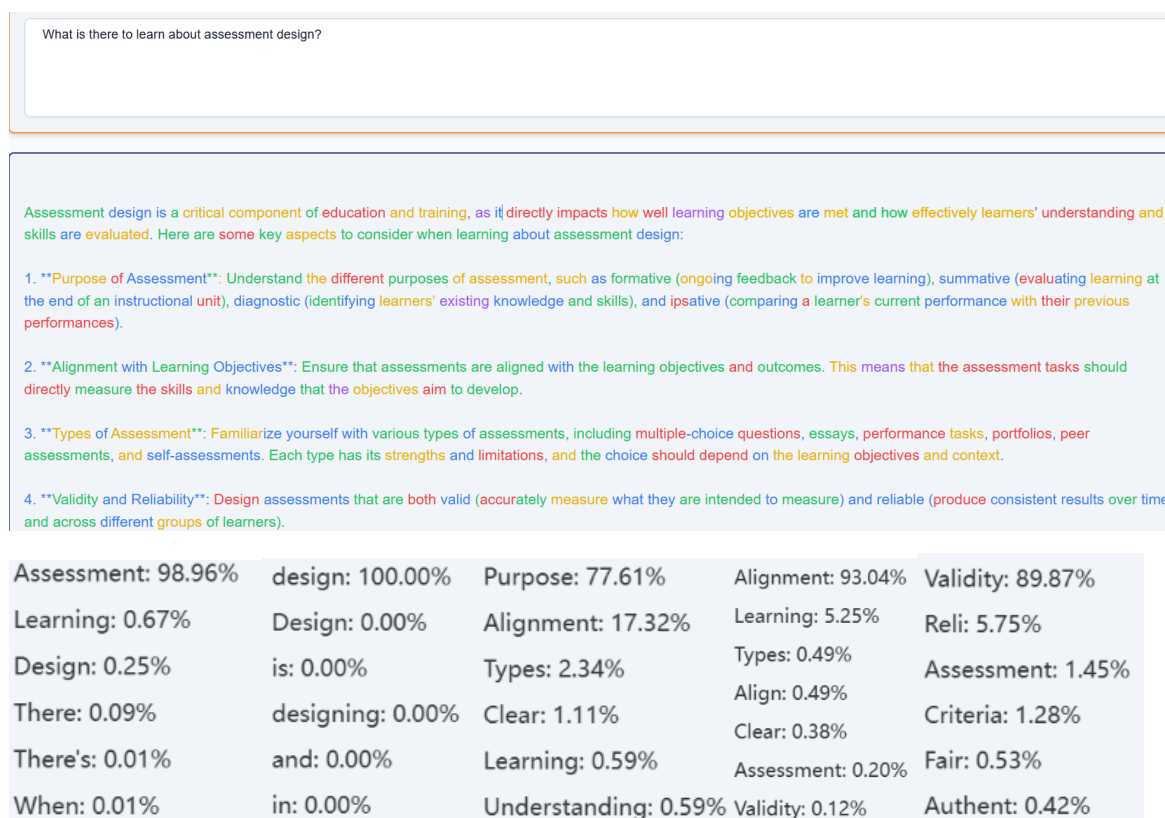
where X is the chosen broad topic. The context window, after using this starter prompt, should contain a list of topics related to X, representing a schema held within the model in its own vocabulary. A user can choose to explore different sub-topics by further prompting using the very tokens the model outputted. Let us use a more detailed example.

Figure 2 shows the response to the following example starter prompt:

What is there to learn about assessment design?

The output has the list of subtopics mentioned earlier with the tokens colour coded, with blue and green representing the higher, and those with red and orange, the lower probabilities. This output is generated using a tool developed by Ryan Tannenbaum called 'tokens for education' (Tannenbaum, 2025), which uses ChatGPT4o. Its temperature was set to 0, which forces it to use the most probable next token from a list of options, as shown for 'assessment' and 'design', which, in that order, have highest probabilities. As is evident here, the starter prompt has helped externalise the model's knowledge structure related to 'assessment design' and the user can delve deeper from this broader output, co-inferring as they continue prompting.

Figure 2. GenAI mind-surfing starter prompt and related future token probabilities.



The follow-up prompt makes use, verbatim, of the externalised knowledge or tokens, to mind-surf and delve deeper, as shown in the second image in Figure 1b. This is analogous to a user clicking links embedded in web pages to surf the web, hence the name mind-

surfing. Using the text 'alignment with learning objectives' from the previous response in the next prompt can serve as a 'link' to generate relevant outputs, by signalling to the attention mechanism what is important to a user, as shown below:

Tell me more about Alignment with Learning Objectives with examples and identify theory that supports this.

Figure 3. GenAI mind-surfing follow-up prompt, delving deeper.

Tell me more about Alignment with Learning Objectives with examples and identify theory that supports this.

Alignment with learning objectives is a foundational principle in assessment design, ensuring that assessments accurately measure the intended learning outcomes. This alignment is crucial for providing meaningful insights into learners' progress and for guiding instructional decisions. Here's a deeper look into this concept, supported by examples and relevant educational theories.

Importance of Alignment

- **Clarity and Focus**: When assessments are aligned with learning objectives, they provide clear targets for both teaching and learning. This helps educators focus their instruction and helps learners understand what is expected of them.
- **Validity**: Alignment ensures that assessments are valid, meaning they measure what they are supposed to measure. This is essential for making accurate judgments about learners' knowledge and skills.

Examples of Alignment

1. **Mathematics**: If a learning objective is for students to "solve quadratic equations using the quadratic formula," an aligned assessment might include problems that require students to apply the quadratic formula to solve equations. An assessment that asks students to simply identify quadratic equations without solving them would not be aligned.
2. **History**: For a learning objective like "analyze the causes and effects of the Industrial Revolution," an aligned assessment could be an essay or project where students explore specific causes and effects, rather than a multiple-choice test that only asks for dates and definitions.
3. **Science**: If the objective is to "understand the process of photosynthesis," an aligned assessment might involve a lab activity where students demonstrate the process or a diagram labeling task, rather than a test that only asks for rote memorization of terms.

Theoretical Support

Several educational theories support the importance of alignment between assessments and learning objectives:

- **Constructive Alignment (John Biggs)**: This theory emphasizes the alignment of teaching activities, learning outcomes, and assessment tasks. According to Biggs, all components of the educational process should be aligned to support deep learning. Constructive alignment ensures that what is taught and assessed is directly related to the learning objectives.
- **Backward Design (Grant Wiggins and Jay McTighe)**: This approach involves starting with the end in mind—identifying desired learning outcomes first and then designing assessments and instructional activities to achieve those outcomes. Backward design ensures that assessments are purposefully aligned with learning objectives from the outset.
- **Bloom's Taxonomy**: This framework categorizes educational goals into levels of complexity and specificity, from basic knowledge recall to higher-order thinking skills like analysis and evaluation. Aligning assessments with the appropriate level of Bloom's Taxonomy ensures that they accurately measure the intended cognitive processes.

The green and blue text (e.g. Examples of alignment) in the response in Figure 3 means highly relevant next tokens, helping to delve deeper and co-inferring, by the user guiding the model to generate highly relevant next tokens and output in the direction they want, usually going from a broad to a narrower focus. The output shows the examples of alignment and theories as were categorically requested by the user. Once you have externalised content related to your topic of interest, you may wish to work with selected

parts of it or bring in other related things that may have been missed out. This is where the next mind-metaphor comes in.

Mind-mending

Mind-mending is the act of exercising user agency in controlling what stays in the context window, either within the output or through new insertions as needed by the user. The human analogy here is when you are guiding someone you are talking with and highlight to them what is more important to consider during your conversation. Controlling the conversation in this way can empower them as they co-infer with the model. The model gets aided in focusing its attention on a selected topic list. Most users are not aware of this when conversing with LLMs. In fact, when using LRMs or agentic AI systems, mind-mending can still be done post mind-surfing and prior to calling a reasoning model or agentic AI system for something specific. Let us see in detail what mind-mending involves. It is not shown in Figure 2, but not all the 10 topics may be of relevance to the user and may contain 'hallucinations' or conjectures. 'Hallucinations' or conjectures are outputs from an LLM that may not meet the needs of the user but are well within the possible range of outputs from a model given what was in the context window. Equally a user may wish to work with some related topics that were missed out in the schema when prompted first, say due to lower probability values. Users could add, remove or correct the schema themselves through mind-mending, as shown in the second image in Figure 1c. When the user attempts to control (or mind-mends) the content in the context window, they control the future token generation in a reliable and desirable way, reducing 'hallucinations' or conjectures, for example, by using the prompt:

I only want to talk about 2, 4, 5 and 6 going forward. and could you add 'assessment security' in the mix.

Thus, mind-mending can and should be used anytime throughout the interactions to control LLM outputs, thereby empowering the user. A new feature in most LLMs, namely Canva, allows such editing of the output. Post mind-mending, the conversation aligns more with the user's way of thinking and working because it is the user who is co-inferring and controlling the LLM output. Figure 4 shows blue coloured (high probability) topic lists to show this control, in response to the mind-mending prompt used. Once you have successfully externalised one schema, you can choose to work with more than one

schema in an expected fashion, or you could choose to be creative to produce novel outputs. The next mind-metaphor takes things one level up into the production of work outputs but also into creativity.

Figure 4. GenAI mind-mending prompt and its control on output.



Mind-bending

Mind-bending is the act of simultaneously attempting to externalise two or more schemata the 'GenAI-mind' or the LLM is trained on to achieve mixing of the two knowledge structures in a meaningful and creative way. Humans can use existing knowledge structures from more than one discipline to produce work or be creative. Here you are

asking the model to use the attention mechanism on multiple different schemata, say one representing a subject area and the other pedagogy. Mind-bending can also be used before the use of LRMs to inform the work accomplished by LRMs. Figure 1d shows how two schemata are used in expected and unexpected ways through mind-bending. For example, an academic may be well trained in using Bloom's taxonomy for designing learning outcomes, just as they are trained within their specialist domain knowledge. Assume that the 'GenAI-mind' can generate outputs about pedagogical theories (Bloom's taxonomy and learning outcomes) and outputs in the user's subject domain. Mind-bending happens when text generation is deliberately influenced by the user using two or more domains of knowledge already present in the model's knowledge structure (or the two schemata in the 'GenAI-mind'). This is achieved by including words that reference these domains together in the same prompt. This way, a user may use GenAI to produce what they already know, or even create new outputs, which may be accurate without them knowing a great deal about them. Mind-bending, therefore, gives the user's agency a productive and/or creative boost (subject to verification of course). In fact, future skills in the orchestration of GenAI tools, working in related fields on demand, to solve complex problems under human control will become important. This could shift the nature of a professional from being T shaped (specialists in one area) (Demirkan and Spohrer, 2015) to being comb shaped (capable of inter, trans, multidisciplinary work or orchestration of the same) (Ninan, Hertogh and Liu, 2022). Mind-bending empowers staff and advanced students to first learn about new domains and then use the knowledge together with other chosen domain(s) they are familiar with already.

Figure 5 shows a mind-bending prompt and the transformation of a single learning objective, 'To analyse the causes and effects of the Industrial Revolution', to multiple learning outcomes written at different levels as an example of this:

Identifying the different levels in Bloom's taxonomy model, map this learning objective, 'To analyse the causes and effects of the Industrial Revolution', to a relevant level and rewrite it to make it suitable to all levels of the model.

Due to the need for verification and selection of established workflows when using this approach, the use of mind-bending is more suitable for subject-matter experts and researchers, but it may also be used by students studying at a higher level. For example, academics could use it to create pedagogically inspired resources, and researchers and

advanced students could use this approach to analyse data using different lenses in the process of producing research work. Of course, they will need to develop and demonstrate their confidence in the tool, using comparative statistics like Kappa (Hua and Yao, 2024). Kappa is used to establish inter-rater reliability between two raters – in this case, it could be between humans and machines – on how each codes research data or makes classification decisions. For an example of this, in the context of scoping reviews, see Malik and Sime, (2025), where we classify synthetic abstracts generated using mind-metaphors. Whilst mind-bending is more useful in doing work within a user's chosen domain or cross-domain creative work, the next few mind-metaphors are about learning with GenAI. These metaphors can help learners engage with the topic in different ways.

Figure 5. GenAI mind-bending prompt and its impact on output.

Identifying the different levels in bloom's taxonomy model, map this learning objective "To analyze the causes and effects of the Industrial Revolution" and rewrite it to make it suitable to all levels of the model.

Bloom's Taxonomy is a hierarchical model used to classify educational learning objectives into levels of complexity and specificity. The original taxonomy, developed by Benjamin Bloom and colleagues, consists of six levels, which were later revised. The revised version includes the following levels:

1. **Remembering**: Recalling facts and basic concepts.
2. **Understanding**: Explaining ideas or concepts.
3. **Applying**: Using information in new situations.
4. **Analyzing**: Drawing connections among ideas.
5. **Evaluating**: Justifying a decision or course of action.
6. **Creating**: Producing new or original work.

The original learning objective, "To analyze the causes and effects of the Industrial Revolution," is situated at the **Analyzing** level. To adapt this objective to all levels of Bloom's Taxonomy, we can rewrite it as follows:

1. **Remembering**: List the key events and figures associated with the Industrial Revolution.
- **Example Task**: Create a timeline of major events during the Industrial Revolution.

2. **Understanding**:
- Objective: "To explain the main concepts and significance of the Industrial Revolution."
- Task: Students will describe the fundamental concepts and explain why the Industrial Revolution was significant.
3. **Applying**:
- Objective: "To illustrate how the Industrial Revolution influenced modern industrial practices."
- Task: Students will use examples to show how the Industrial Revolution's innovations are applied in today's industries.
4. **Analyzing**:
- Objective: "To analyze the causes and effects of the Industrial Revolution."
- Task: Students will examine and break down the causes and effects, identifying relationships and patterns.
5. **Evaluating**:
- Objective: "To assess the impact of the Industrial Revolution on different social classes."
- Task: Students will evaluate the positive and negative impacts of the Industrial Revolution on various social groups, providing evidence to support their judgments.
6. **Creating**:
- Objective: "To design a presentation proposing an alternative historical scenario where the Industrial Revolution occurred in a different region."
- Task: Students will create a presentation that imagines how the Industrial Revolution might have unfolded differently if it had started in another part of the world, considering potential causes and effects.

Mind-stretching

Mind-stretching empowers a learner, researcher or staff member to explore concepts using two specific or random axes to help them think clearly and completely about something or to stretch their own thinking. There are many human analogies here, for example Boston Consultancy Group’s Growth Share Matrix. This metaphor is best used with a foundational LLM but could be used post LRM use to analyse the output as well. Mind-stretching is an extension of the mind-bending technique, but for finer analysis as opposed to producing text. It is akin to zooming in on a topic with multiple lenses, that is, it requires the LLM to view a topic using two scaled axes, forming a two-by-two matrix. For example, the following prompt produces Table 2, which shows how the ‘GenAI-mind’ categorises various entries in the table into four different quadrants, followed by some explanation of the logic behind this (not shown):

Suggest a two-by-two matrix to better understand the industrial revolution.

Any two specific lenses of your choice could be specified instead, enhancing the agency here. Students could use it to understand and be more critical about a topic, whilst researchers could use this approach to analyse new possibilities with ease. The skill in using this technique relies on choosing meaningful axes and interpreting the results. Tapping into the ability of GenAI to play multiple roles in a simulated conversation to learn new topics is the focus of the mind-metaphors described next.

Table 2. Mind-stretching helps in the analysis of a concept using multiple lenses (generated by ChatGPT 4.5).

Industrial Revolution Matrix

	Low Technological Advancement	High Technological Advancement
Negative Societal Impact	Pre-Industrial Society - Agricultural economy - Manual labor dominates - Limited social mobility	Early Industrial Revolution - Harsh working conditions - Child labor & pollution - Displacement of artisans
Positive Societal Impact	Sustainable Traditionalism - Craft-based economies - Small-scale innovations - Stability in social structures	Modern Industrial Economy - Technological progress - Improved labor laws - Increased living standards

Mind-storming

Mind-storming is the act of generating an appropriate simulated discussion between multiple role holders, for example a teacher and a small group of students, to empower students to learn from the simulated discussion. So far, the mind-metaphors focused on controlled content creation. Next, I show how mind-storming triggers a simulated discussion, to play out on a users' screen, between chosen virtual role players. For example, between an excellent virtual professor and their novice and advanced students. The virtual professor's useful questions, the novice students' common mistakes and the support from other more advanced virtual students could all be defined as part of the system prompt within a GenAI tutor. A system prompt is a prompt which is appended into the context window by the tutor designer, hidden from the student. This empowers the designer to set up a pedagogically helpful interaction for their students to benefit from. The output of such a prompt could help real students to learn from the simulated interaction on their screens. When used as a simulated discussion between virtual key role holders, working within an institution or between them and external accreditors, it could help prepare for an accreditation visit, tackling challenges and problem solving. Researchers could use this method to prepare for data gathering in a simulated focus group and analysis. This approach lends itself well to learning by reading the contrasting conversation between virtual roles. The next mind-metaphor turns this simulated conversation into a hybrid of virtual and real conversation for a more realistic and engaging feel.

Mind-gaming

Mind-gaming is the act of taking part in a challenging real-time simulated discussion between multiple virtual role holders and the user. For example, a mind-gaming session with a small group of simulated students with pre-defined roles and a challenging brief, can empower an early career academic to learn how to respond to students in a challenging situation. Mind-gaming is a simulated, but more challenging, conversation than mind-storming. It could be used to prepare a PhD candidate for the panel viva, for example. Researchers could use this as a method to practice netnography in virtual settings described in the scenario and analyse the data collected themselves. Academics and students could use it for group interviews to prepare for job seeking.

Challenges and barriers

GenAI use in higher education still faces several structural and epistemic barriers which may delay an empowered emergence from the said liminal space. Some unresolved questions, related to data privacy, intellectual-property ownership and sustainable use, do create barriers for staff and students as well as institutional risk that need understanding, demystifying and/or mitigating. Staff and students should not become over-reliant on GenAI and should exercise their own thinking, using the mind-metaphors, so that they can reduce the 'hallucinations' or produce more relevant and accurate outputs. They need to be aware of the limitations in training data by exploring the model's knowledge structures through mind-surfing, and control what is in the context window through mind-mending. Bias and opacity of models is another problem that undermines academic trust from staff and students. Compute-intensive inference (Snell et al., 2024) and the associated energy footprint more generally, may conflict with the universities' sustainability targets, so understanding the impact of comparative human work may become critical. The benefits of co-inferring and collaboration can only be realised through sustained professional development and with inclusive and sustainable design.

Conclusion

This paper has highlighted the differences between foundational and other models in terms of their impact on human and machine agency. It has highlighted how mind-metaphors can empower staff and students to better understand and master GenAI literacy, to help them emerge out of the liminal space created by GenAI developments. Staff and students can use mind-metaphors with foundational LLMs or before and after the use of LRMs to direct the attention mechanism, so the models generate outputs that are less 'hallucinatory' and more aligned with the user's intent. With mind-metaphors the user can control what is in the context window to influence the predicted output token and direct the model's attention mechanism. Staff and students need to develop their agency in the use of foundational LLMs to emerge empowered out of the liminal space. LRMs and agentic AI systems are dependent on the context window and attention mechanism too. Although LRMs and agentic AI systems take agency away from users, careful mind-surfing and mind-mending can help control the input to LRMs and agentic GenAI systems.

By focusing on the activities, such as mind-surfing, mind-mending and mind-bending, users are empowered to talk about their use of the GenAI in a way that promotes their agency over it as opposed to referring GenAI as a noun, for example 'GenAI can do...', which suggests the agency is with the models. The framework has helped over 250 academic staff to use foundational models as evidenced by their high satisfaction levels (Malik, 2025). Many lecturers use the techniques to co-infer and collaboratively create content at my institution. I have used these to produce course materials and synthetic abstracts for a GenAI-assisted scoping review screening, with high accuracy and Kappa values when compared with human screening (Malik and Sime, 2025).

Students, too, can start by mind-surfing and mind-mending their chosen domain and/or metacognitive topics. This can be done in supported classrooms where they are encouraged to verify LLM outputs as they learn with GenAI. They can learn to be critical in this way. They can be guided by staff on useful pedagogical concepts such as common misconceptions, troublesome and threshold concepts, which can then be used in mind-storming and mind-gaming to learn and practise concepts. Advanced students can co-infer and collaborate by using mind-stretching and mind-bending to support their study and work needs. The framework can be used in many other use cases, within teaching, learning, research and beyond, and I invite researchers and practitioners to report on such use cases in appropriate conferences and journals in future.

Acknowledgements

The author used the following generative AI tools in the preparation of this manuscript: ChatGPT. The tasks performed by ChatGPT include formatting of references exported from Google Scholar and image generation of Figure 1 (a) – (d). The author takes full responsibility for the content of the article and has complied with the JLDHE's principles of AI use.

References

Abels, C. M., Lopez-Lopez, E., Burton, J.W., Holford, D.L., Brinkmann, L., Herzog, S.M. and Lewandowsky, S. (2025) 'The governance & behavioral challenges of

generative artificial intelligence's hypercustomization capabilities', *Behavioral Science & Policy*, 11(1), pp.22-32. Available at:
<https://doi.org/10.1177/23794607251347020>

Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J. et al. (2023) 'Gpt-4 technical report', *arXiv preprint arXiv:2303.08774*. Available at: <https://doi.org/10.1177/23794607251347020>
(Accessed: 30 November 2025).

Akbulut, C., Weidinger, L., Manzini, A., Gabriel, I. and Rieser, V. (2024) 'All too human? Mapping and mitigating the risk from anthropomorphic AI', *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 7(1), pp.13-26. Available at:
<https://doi.org/10.1609/aies.v7i1.31613>

Andersen, J.P., Degn, L., Fishberg, R., Graversen, E.K., Horbach, S.P., Schmidt, E.K., Schneider, J.W. and Sørensen, M.P. (2025) 'Generative artificial intelligence (GenAI) in the research process – a survey of researchers' practices and perceptions', *Technology in Society*, 81, p. 102813. Available at:
<https://doi.org/10.1016/j.techsoc.2025.102813>

Cascella, M., Semeraro, F., Montomoli, J., Bellini, V., Piazza, O. and Bignami, E. (2024) 'The breakthrough of large language models release for medical applications: 1-year timeline and perspectives', *Journal of Medical Systems*, 48, p. 22. Available at:
<https://doi.org/10.1007/s10916-024-02045-3>

Concannon, F., Costello, E., Farrell, O., Farrelly, T. and Wolf, L.G. (2023) 'There's an AI for that: rhetoric, reality, and reflections on EdTech in the dawn of GenAI', *Irish Journal of Technology Enhanced Learning*, 7(1). Available at:
<https://doi.org/10.22554/ijtel.v7i1.116>

Coulardeau, J. (2018) 'No cognition without communication', *Journal of Psychology Research*, 8(9), pp. 443–477. Available at:
<https://api.semanticscholar.org/CorpusID:149655973> (Accessed: 30 November 2025).

- Demirkan, H. and Spohrer, J. (2015) 'T-shaped innovators: identifying the right talent to support service innovation', *Research-Technology Management*, 58(5), pp.12–15. Available at: <https://doi.org/10.5437/08956308X5805007>
- Dubey, S.S., Astvansh, V. and Kopalle, P.K. (2024) 'Generative AI solutions to empower financial firms', *Journal of Public Policy & Marketing*. Available at: <https://dx.doi.org/10.2139/ssrn.5050951>
- Faruqui, N., Raju, N.V.D.S.S.V.P., Sivakumar, S., Patel, N., Bhaskaran, S.V., Khanam, S. and Bhuiyan, T. (2025) 'Gen-Optimizer: a generative AI framework for strategic business cost optimization', *Computers*, 14(2), p.59. Available at: <https://doi.org/10.3390/computers14020059>
- Freeman, J. (2025) 'HEPI/Kortext AI survey shows explosive increase in the use of generative AI tools by students', *Higher Education Policy Institute (HEPI)*, 26 February. Available at: <https://www.hepi.ac.uk/2025/02/26/hepi-kortext-ai-survey-shows-explosive-increase-in-the-use-of-generative-ai-tools-by-students/> (Accessed: 5 March 2025).
- Hua, H. and Yao, C.J. (2024) 'Investigating generative AI models and detection techniques: impacts of tokenization and dataset size on identification of AI-generated text', *Frontiers in Artificial Intelligence*, 7, article 1469197. <https://doi.org/10.3389/frai.2024.1469197>
- Hua, W. and Zhang, Y. (2022) 'System 1+ system 2= better world: neural-symbolic chain of logic reasoning', *Findings of the Association for Computational Linguistics: EMNLP*, Abu Dhabi December, pp.601–612. Available at: <https://doi.org/10.18653/v1/2022.findings-emnlp.42>
- Kahneman, D. (2011) *Thinking, fast and slow*. Allen Lane and Penguin Books, New York.
- Kosmyna, N., Hauptmann, E., Yuan, Y.T., Situ, J., Liao, X.H., Beresnitzky, A.V., Braunstein, I. and Maes, P. (2025) 'Your brain on ChatGPT: accumulation of cognitive debt when using an AI assistant for essay writing task', *arXiv preprint arXiv:2506.08872*. Available at: <https://doi.org/10.48550/arXiv.2506.08872>

- Lakoff, G. and Johnson, M. (2008) *Metaphors we live by*. Chicago: University of Chicago Press.
- Lee, H.P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R. and Wilson, N. (2025) 'The impact of generative AI on critical thinking: self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers', *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Yokohama, Japan 26 April-1 May. Available at: <https://doi.org/10.1145/3706598.3713778>
- Li, S., Liu, J. and Dong, Q. (2025) 'Generative artificial intelligence-supported programming education: effects on learning performance, self-efficacy and processes', *Australasian Journal of Educational Technology*, 41(3), pp.1–25. Available at: <https://doi.org/10.14742/ajet.9932>
- Malik, M. (2025) *Mind-metaphors for co-intelligence and collaboration with GenAI for teaching and learning*. YouTube. Available at: <https://youtu.be/SZj8qittRjQ?si=5ZzCHeKENzE3Ad17> (Accessed: 30 November 2025).
- Malik, M. and Sime, J.A. (2025) 'Teamwork, co-regulation, and socially shared regulation skills within Engineering education studies: a GenAI-assisted scoping review', *2025 ASEE Annual Conference & Exposition*. Montreal, Quebec, Canada 22-25 June. Available at: <https://peer.asee.org/57199> (Accessed: 30 November 2025).
- Nikolic, S., Wentworth, I., Sheridan, L., Moss, S., Duursma, E., Jones, R.A., Montserrat Ros, M. and Middleton, R. (2024) 'A systematic literature review of attitudes, intentions and behaviours of teaching academics pertaining to AI and generative AI (GenAI) in higher education: an analysis of GenAI adoption using the UTAUT framework', *Australasian Journal of Educational Technology*, 40(6), pp.56–75. Available at: <https://doi.org/10.14742/ajet.9643>
- Ninan, J., Hertogh, M. and Liu, Y. (2022) 'Educating engineers of the future: T-shaped professionals for managing infrastructure projects', *Project Leadership and Society*, 3, article 100071. Available at: <https://doi.org/10.1016/j.plas.2022.100071>

- Oakley, B., Johnston, M., Chen, K., Jung, E. and Sejnowski, T. (2025) 'The memory paradox: why our brains need knowledge in an age of AI', in *The future of artificial intelligence: economics, society, risks and global policy* (Springer Nature, forthcoming). Available at <https://dx.doi.org/10.2139/ssrn.5250447>
- Opus, C. and Lawsen, A. (2025) 'Comment on the illusion of thinking: understanding the strengths and limitations of reasoning models via the lens of problem complexity', *arXiv preprint arXiv:2506.09250*. Available at: <https://arxiv.org/abs/2506.09250>
- Roberts, A.I. and Roberts, S.G. (2020) 'Communicative roots of complex sociality and cognition', *Biological Reviews*, 95(1), pp.51–73. Available at: <https://doi.org/10.1111/brv.12553>
- Senoner, J., Schallmoser, S., Kratzwald, B., Feuerriegel, S. and Netland, T. (2024) 'Explainable AI improves task performance in human–AI collaboration', *Scientific Reports*, 14(1), 31150. Available at: <https://doi.org/10.1038/s41598-024-82501-9>
- Shojaee, P., Mirzadeh, I., Alizadeh, K., Horton, M., Bengio, S. and Farajtabar, M. (2025) 'The illusion of thinking: understanding the strengths and limitations of reasoning models via the lens of problem complexity', *arXiv preprint arXiv:2506.06941*. Available at: <https://doi.org/10.48550/arXiv.2506.06941>
- Shardlow, M. and Przybyła, P. (2024) 'Deanthropomorphising NLP: can a language model be conscious?', *PLoS One*, 19(12), article e0307521. Available at: <https://doi.org/10.1371/journal.pone.0307521>
- Singh, R.G. and Ngai, C.S.B. (2024) 'Top-ranked US and UK's universities' first responses to GenAI: key themes, emotions, and pedagogical implications for teaching and learning', *Discover Education*, 3(1), article 115. Available at: <https://doi.org/10.1007/s44217-024-00211-w>
- Snell, C., Lee, J., Xu, K. and Kumar, A. (2024) 'Scaling LLM test-time compute optimally can be more effective than scaling model parameters', *arXiv preprint*

arXiv:2408.03314. Available at: <https://arxiv.org/abs/2408.03314> (Accessed: 30 November 2025).

Sperber, D. and Wilson, D. (1987) 'Précis of relevance: communication and cognition', *Behavioral and Brain Sciences*, 10(4), pp. 697–710. Available at: <https://doi.org/10.1017/S0140525X00055345>

Tannenbaum, R., (2025) *Tokens for education*. Available at: <http://tokens.for.education/> (Accessed: 30 November 2025).

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L. and Polosukhin, I. (2017) 'Attention is all you need', *31st Conference on Neural Information Processing Systems (NIPS 2017)*. Long Beach, CA, USA. Available at: <https://arxiv.org/pdf/1706.03762>

Author details

Manish Malik is Associate Dean (Education) for the School of Science, Psychology, Arts and Humanities, Computing, Engineering and Sports at Canterbury Christ Church University, leading academic, quality, curriculum, assessment, and student success. His academic and research interests include AI in education, engineering education, educational technology, assessment enhancement and student experience.

Licence

©2025 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <http://creativecommons.org/licenses/by/4.0/>. Journal of Learning Development in Higher Education (JLDHE) is a peer-reviewed open access journal published by the Association for Learning Development in Higher Education (ALDinHE).