Abstract

Frictionless, dematerialised environments (FDEs) are producing a paradox: a rise in cognitive exhaustion and somatic discomfort even as information environments become objectively more efficient. This editorial argues the progressive removal of material, bodily, and social resistance from knowledge environments does not expand cognitive possibility; it forecloses the conditions under which understanding has historically been assembled. What distinguishes the current moment is not merely the pace of change but the nature of the niche: for the first time, the dominant cognitive environment is genuinely immaterial and acquiescent, aligned with the Cartesian fantasy of mind liberated from body rather than with the embodied, situated reality of human cognition. It introduces the concept of the cognitive uncanny valley to account for the somatic discomfort that accompanies AI-generated knowledge. The editorial concludes with a call not for greater cognitive efficiency but for its opposite: the deliberate preservation of inefficiency, impasse, and not-knowing.

The Effort-Reward Mismatch: Cognitive Friction and the Conditions of Possibility

This is the most information-rich environment in human history (Hołyst et al., 2024); an environment marked not only by information abundance but also ease of access and concentration in digital interfaces—what Bawden and Robinson refer to as "homogenised diversity" (2009, p. 181). These frictionless, dematerialised environments (FDEs) are characterised by instant access to information while the effort, materiality, and embodied engagement that have historically accompanied knowledge acquisition have been progressively designed out. In these environments, information is available at people's fingertips with little or even no effort. Despite this, there are many reports of people feeling cognitively exhausted and overwhelmed even when their attitudes are generally positive and their thinking objectively more efficient (Shalu et al., 2026). As Constance de Saint Laurent (2024) has clearly articulated in an earlier editorial in this journal, technology both expands and contracts possibilities and often for different groups of people.

The arrival of large language models (LLMs) has accelerated this environmental shift to a qualitatively new level. This is not merely a shift in scale, but in shift in kind: a technology performs the synthesis, evaluation, and articulation of knowledge on the agent's behalf, leaving the human as passive recipient rather than active thinker. We are all plugged into the Matrix and we no longer even need to know Kung Fu, the machines do the fighting for us! A growing body of empirical research documents the cognitive consequences of this LLM and human coupling. Famously, Kosmyna et al. (2025) used electroencephalography to monitor brain activity across four sessions of AI-assisted essay writing, finding that LLM users exhibited substantially reduced neural connectivity compared with those writing without tools, that 83% could not quote from the essay they had just written, and that those who had relied on AI struggled to recover independent cognitive engagement when the tool was removed. Gerlich (2025) found a significant negative correlation between frequent AI tool usage and critical thinking abilities. Stadler et al. (2024) showed that LLMs reduce cognitive load but compromise depth of scientific enquiry. Lee et al. (2025) found self-reported reductions in cognitive effort and decision-making confidence among knowledge workers using generative AI regularly. The list grows rapidly and exponentially[1].

[1] Of course there is a scientific story here. The finding that cognitive functions are retained when coupled with LLMs fits less into the current zeitgeist and so those papers may find it harder to see the light of day. There is an active and reciprocal relationship between publication bias and our understanding of psychological phenomena that actively constructs our understanding.

On the surface to proponents of extended cognition, this coupling should not be so concerning. Embodied and situated human agents have always structured the world to reduce cognitive load so the relationship between offloading onto digital tools and cognitive ease is to be expected (Kirsh, 1995; Vallée-Tourangeau, 2013; Wilson, 2001). Even before the current information explosion, there has been too much information for the human brain to comfortably process. Theories of cognitive extension (such as the Extended Mind Theory; Clark & Chalmers, 1998) have long documented the offloading of cognitive work to the environment and increasing capacity in this way. Humans have a history of manipulating the environment to construct cognitive niches where the environment both scaffolds and becomes an active agent in their cognition (Clark, 2008; Malafouris, 2013; Sterelny, 2012). However, the difference here lies in the somatic discomfort that people experience when coupled over the long term with digital tools in FDEs leading to reported brain fog and "AI brain fry" (Bedard, J. et al., 2026). This difference is important to note because not all forms of cognitive offloading to the digital world produce this *discomfort,* even when they reduce the capacity of the individual to work independently. For example, there has been concern over the effect of Google on the tendency to understand where to find information rather than committing that information to memory (Gong & Yang, 2024). However, in an environment in which the agent has learnt that the information is stored in stable place, knowing who knows what and knowing what you do not know are sophisticated and genuinely adaptive cognitive skills (see the argument on "glue and trust" in the original Extended Mind Thesis; Clark & Chalmers, 1998). From this perspective, phenomena such as the so-called Google Effect do not represent cognitive failure but a rational adaptation to a changed information environment. However, in the case of the current human-technological coupling there are embodied indications that this environmental shift is different.

Of course, it is possible that our brains and behaviours will adapt to these FDEs over time and it is simply the rapid accumulation of change that has caused these somatic side effects rather than anything deeper. Analogies are often drawn between this moment of information transition and other moments such as the widespread adoption of the calculator (Anderson et al., 2025). From this view, the somatic discomfort that accompanies the current change is a reflection by how rapidly it has spread rather than anything more serious. On the other hand, it is also possible that what distinguishes the current moment from previous episodes of cognitive niche construction is not only the pace of change but also the nature of the cognitive niche being constructed (see Voinea et al., 2026 for differences between calculators and LLMs). Previous cognitive niches have been socio-material in nature and

while malleable retained their own form independent and distinct from the individual human mind; abstraction was supported by the concrete. What is new about the niche that we see in FDEs is that it is genuinely immaterial and, importantly, acquiescent. These environments are based on pure language and on disembodied symbol where there are no constraints and any fantastical imagining is possible. It is, in other words, the first cognitive niche that fundamentally resembles the Enlightenment fantasy of an intellect not restrained by the inflexibility of the material world or the weakness and unreliability of the body. That is precisely what makes it at once so seductive, and yet, I argue, so poorly matched to our actual cognitive profile.

Fundamentally, the current situation of cognition coupled with FDEs reflects the dominant perspective running through Western thought from Plato to Descartes to the transhumanists: that the mind is the part of us that transcends the body, that knowledge outlives flesh and that intelligence properly understood is liberated from the constraints of a physical self that can only age, degrade, and die. AI did not invent this perspective; it is merely its most technologically persuasive expression yet. From this view, effortless, disembodied information acquisition fits perfectly with the promised freedom of mind detached from body and world (Ross et al., 2024). However, the negative somatic effects that are reported point to a dissonance between the promise of Enlightenment Cartesian dualism and the embodied thinking that the situated agent actually does. This discontent can be explained by something fundamental; there is an ongoing reconfiguration of cognitive life in the model of rational, abstracted thought, which is poorly matched to the thinking of embodied, situated, socially assembled beings. The emerging evidence does not show a mind being damaged so much as a mind being reshaped in ways that are increasingly misaligned with the conditions under which human cognition has historically flourished.

Possibility Studies offers a direct challenge to this view in the ongoing commitment that the possible should not be romanticised. The belief that possibility expands without limit when constraints are removed, enforced by cultures oriented towards progress and technological liberation, can itself become a source of cognitive and psychological harm when the expectations it generates are not met. In addition, constraints are not the enemies of possibility but its enablers, and encountering resistance can itself be a trigger for genuine discovery (Glăveanu, 2023). From this perspective, FDEs do not expand possibility even though on the surface they make anything possible. In promising to remove all friction, they foreclose the very conditions under which genuine possibility emerges, whilst simultaneously

generating the dissonance and discomfort that people report when coupled with them over time.

Effort, Scarcity, and the Evolutionary Roots of Possibility

To understand this reconfiguration, it helps to borrow an analogy from another domain. The industrialisation of food production in the nineteenth and twentieth centuries created an environment of engineered caloric abundance (Bellisari, 2008; Gluckman, 2006). In this environment, the reward signals that evolution had calibrated for scarcity, such as the powerful drive towards fat, sugar, and salt that had reliably guided our ancestors towards nutritionally valuable food in conditions where such food was rare, became maladaptive when they were not accompanied by the physical work that had always accompanied them. The result was a systematic mismatch between a reward system shaped by one kind of environment and a world that had changed faster than biology could follow. The analogy is not perfect, and there are dangers in overextending it, but the structural logic is similar to what we are now observing in the cognitive domain: a system calibrated for one set of conditions, operating in conditions for which it was not designed, producing a form of life that feels satisfying in the moment but is increasingly misaligned with what the system needs to function well. From the perspective of our argument, both environments now provide resources without the effort that we evolved to exert.[2]

Human cognition evolved in conditions of information scarcity. The evolutionary mismatch hypothesis offers a useful starting point for understanding the new cognitive environment. Li et al., (2018) argue that human psychological mechanisms are adaptations that evolved to process particular kinds of environmental inputs, and that when those inputs change faster than the mechanisms can follow, the result is a system producing outputs it was not shaped for. While Li and colleagues do not consider cognitive adaptions, Pirolli and Card's theory of Information Foraging (1999) demonstrates that human agents approach

[2] It is worth noting that the parallel runs deep. Long before the industrialisation of food production, the expansion of the sugar trade in the sixteenth and seventeenth centuries introduced refined sweetness into diets at a scale and concentration that had no precedent in human evolutionary history. But the sugar trade also tells a darker story: its abundance was made possible by the violent exploitation of people in the service of satisfying appetites elsewhere. Transitions of this kind are never merely technological or economic, rather they reorganise who bears the cost and who receives the benefit, and the cognitive transition we are describing is unlikely to be any different. The communities most exploited in the sugar and slave trades are disproportionately represented among those performing the poorly paid, psychologically damaging data labelling and content moderation work on which contemporary AI systems depend (see Perrigo, 2023; Hao, 2023). The costs of cognitive abundance, like the costs of caloric abundance before it, are not borne by those who most benefit from it.

information seeking the way animals approach food: assessing the likely yield of any given source against the effort required to exploit it and allocating attention accordingly. They suggest that human agents are, in a formulation borrowed from Dennett, *informavores*, that is creatures whose cognitive architecture is built around the active, effortful pursuit of knowledge. Even theories of information encountering rather than active search that arose as information became more able to be passively acquired emphasise the active role of the person in making the most of the information that is encountered (Erdelez & Makri, 2020). We are primed to explore for information. From an evolutionary perspective, it used to require sustained investment to acquire reliable and trustworthy knowledge; knowledge was acquired through observation, social learning, trial and error and importantly physical engagement with the world and with other people. Information acquired through this sort of effort was more likely to have been tested against reality and also more likely to be encoded deeply and integrated with prior knowledge. This allowed it to be consolidated into the kind of understanding that can be retrieved, transferred, and built upon.

Those impoverished information environments have in turn shaped emotional geographies, particularly the importance of epistemic emotions to support and motivate information search. For example, we can draw on Loewenstein's (1994) Information Gap Theory of curiosity, in which he argues that curiosity is the awareness of a gap between what we know and what we want to know. What is important to note is that this model of curiosity does not predict mere interest but epistemic discomfort, which then translates to a motivational state that drives active search. Crucially, the dopaminergic reward circuitry is activated by the effort of closing the gap, not merely to the arrival of information (Gruber et al., 2014). This suggests that the act of foraging itself is important. In work that we are currently doing in our lab, my colleague Selene Arfini and I gave participants hard problems to solve and allowed them to opt to be given the answer rather than go on to solve. Only 17% opted to have the answer prior to offering a solution, the rest wanted to attempt the problem even though there was no financial nor clear epistemic benefit. The satisfaction of solving a problem without help and the sense of ownership over hard-won knowledge are the fundamental motivational drivers through which effortful engagement is sustained and through which knowledge seekers retain a sense of agency. The answer to our problems may come more quickly and efficiently with coupling with FDEs but this is only truly optimal in a model which values the outcome over the process. Possibility Studies frames this affective dimension of cognition as central rather than incidental. Experiences of the possible are multifaceted: they involve the entire being and carry a strong motivational and emotional

dynamic (Glăveanu, 2023). These are not merely cognitive events but possibility experiences, and their affective texture is part of what makes them. Removing the need for search flattens the affective landscape within which genuine possibility is encountered and pursued.

Cognitive Load, Serendipity, and the Instability of Efficiency

Sweller's (1988) Cognitive Load Theory offers a further lens through which to understand what is at stake when we remove all effort from knowledge acquisition. The theory distinguishes between extraneous cognitive load and germane load. Extraneous cognitive load represents effort spent on unnecessary complexity that impedes learning, which Sweller argues should be reduced, while germane load is that the effortful engagement with cognitive material that drives the formation of schemas and the consolidation of understanding. The outsourcing of information synthesis to FDEs does not distinguish between extraneous and germane load, by reducing one, it reduces both. As Chirayath et al. (2025) observe, technology that reduces germane load has a negative impact on deeper engagement, potentially alleviating present burdens whilst risking the erosion of long-term capacity. There is an ongoing cognitive value to the effortful processing of information.

However, even the distinction between extraneous and germane load is less stable than it first appears when we consider cognition as a process extending over time. The distinction as originally formed is essentially synchronic: it evaluates cognitive effort at the moment of processing and classifies friction as either productive or obstructive on the basis of its immediate relationship to the learning goal. However, outcomes from embedded and situated cognitive tasks can take place over several different timescales. What registers as wasted effort or irrelevant noise in the moment frequently turns out, in retrospect, to have been the generative detour that mattered most. The value of these impasses become visible only later, and only because they were not smoothed away. To apply Sweller's distinction in real time is to make exactly the kind of advance determination about which friction is productive that no external agent is in a position to make.

This temporal problem deepens when we consider what FDEs implicitly claim to do. They are designed on the assumption that cognitive friction is, in general, extraneous: an obstacle between the agent and the information they seek, to be minimised in the interest of efficiency. But this is not a neutral design decision. It is an epistemological commitment based on the assumption that the path between question and answer is best shortened, that the detour is waste, that the period of impasse is of less interest than the insight which follows. Sweller's own framework offers no mechanism by which an external system could reliably

distinguish, at the moment of encounter, which difficulty will prove generative and which will not. This is a structural decision and commitment to a view of knowledge that values outcome over process. The other position views the notion of load as a dynamic interaction between person and material. The resistance that would be extraneous noise to an experienced practitioner may be precisely the generative challenge a novice requires; the difficulty that one person navigates towards insight another encounters as mere obstruction. The category is not fixed in advance.

Possibility Studies describes this as the intertwining of the possible and the actual: possibility and actuality continuously feed into and transform each other through the course of action, and it is the tension between them that generates discovery. The margins of the story of scientific discovery are covered with examples of progress made through the encounter with failure and accidence (Copeland, 2019; Ross, 2022; Yaqub, 2018) Serendipitous encounters are not departures from reality but collisions with it. FDEs, in smoothing away that tension, do not open up the possible. They close it down. Failed attempts act as cognitive markers, reorienting the mind towards what is actually important by disrupting a planned cognitive trajectory and forcing a restructuring of the problem space (Ross, 2023b, 2024; Seifert, 2024; Seifert et al., 1994). The external world intrudes, derails, and in doing so opens possibilities that the original trajectory would never have reached. Selene Arfini (2019, 2023) deepens this in her work on ignorant cognition, arguing that the assumption that searching and finding are tightly and linearly coupled is itself a philosophical mistake and that understanding knowledge acquisition requires taking seriously the generative role of not-knowing in the cognitive economy.

The Cognitive Uncanny Valley and What the Data Cannot Capture

It is important to note that the mismatch that is occurring is not about a separation or tension between the mind and the world. This would continue to uphold the misplaced dualistic notion that undergirds much of current thought. Rather, from a perspective of metaplasticity (Malafouris, 2015), what matters is not that there is some original or optimal form of mind to which we should return but rather the continuity of the conditions under which cognitive becoming has historically occurred. The mismatch in question is between a historically constituted process of cognitive becoming and an environment that is systematically removing the conditions that process requires. The mode of cognition where information was assembled through sustained engagement with resistant materials, with the body's own effort, and with the unpredictable resistance of other people, is being displaced by

a mode that is systematically immaterial and non-resistant. This is not just in one part of our cognitive ecologies, it is rapidly expanding to encompass it all. One of these modes is continuous with the embodied, material, socially situated history of human cognition, and one aligns with the dominant theory of abstracted thought and effortless knowledge gain.

This reframing sends us back to the studies that are examining the cognitive impact of coupling with the FDEs. These measure cognitive debt, critical thinking scores, and memory retention which are all metrics that presuppose the brain is essentially an information processing machine to be evaluated in terms of efficiency and output. They ask whether FDEs are making human agents better or worse at cognitive production with the implicit assumption that worse and better can be understood in terms of cognitive efficiency. But this is precisely the model of mind that this editorial has been arguing against. In framing the question this way, these studies risk validating the very framework that produced the problem. Instead, it is possible to recruit other ways of knowing that are not easily measured (De Jaegher, 2019). For example, there is a widespread and growing need to know whether something was written by a human or generated by AI, even when people struggle to say exactly *why* it matters. After all, the intellectual argument that it should not matter is perfectly coherent with a model that is only interested in the answer. This discomfort is an important bodily signal. Something is being registered below the threshold of conscious analysis, a somatic discomfort that precedes and exceeds what we can articulate intellectually.

This could be seen as the cognitive uncanny valley. The concept borrows structural logic from Mori's (1970; 2012) original uncanny valley in robotics: as a robot's human likeness increases, our affinity for it rises, until it reaches a point of near-human but detectably imperfect resemblance, at which point our response abruptly shifts to eeriness and revulsion. The discomfort is not caused by the thing being too different. It is caused by it being almost right, close enough to trigger our recognition systems, but wrong in ways the body detects before the intellect can name them. The structural logic that is transferable here is this: a stimulus activates recognition responses it then fails to satisfy, and the mismatch between the triggered response and the absent confirmation produces discomfort rather than neutrality. In our case, the trigger is not visual human-likeness but epistemic authority. AI-generated text has the form and the apparent authority of knowledge. But something is absent, and our embodied cognitive system registers this absence as wrongness even if our rational selves cannot explain it.

David Pye's distinction between the workmanship of risk and the workmanship of certainty may also be helpful here (Pye, 1968). Originally formulated to understand the

different between craft and factory-based work, the workmanship of risk is work in which the quality of the outcome is not predetermined. It characterises work that depends throughout on the judgement of the maker and in which something is at stake in the doing of it (Ross & Groves, 2023). The workmanship of certainty is work in which the result is predetermined before a single thing is made. Pye's observation was that all the works humans have most admired throughout history have been made by the workmanship of risk. It is possible to somehow recognise the effort and the possibility of failure, and when that is absent, when what we receive is the workmanship of certainty dressed in the appearance of the workmanship of risk, we feel the uncanniness even before we can say what we are feeling.

It is important to establish that this response is not nostalgia, and not merely a social attitude towards AI that will shift as the technology becomes familiar. This may be the case and it is not a given that the change will be negative. The argument is that we should grant the discomfort functional rather than merely sentimental status. This claim carries an empirical prediction: if the cognitive uncanny valley is a genuine perceptual signal, the discomfort it produces should track something real. It should, for instance, correlate with the ability to detect the absence of epistemic markers even when participants cannot consciously articulate what they are responding to. A response that is purely attitudinal or nostalgic would show no such correlation. It is equally important to note that not all forms of cognitive offloading produce this form of discomfort. The discomfort people are describing is triggered by something more specific: the increasing smoothness of our information environments and the progressive removal of every trace of effort, risk, and human struggle from the knowledge we receive. There is a collective *feeling* something is wrong which cannot always be articulated. The dominant intellectual and cultural framework, which relies on the model that efficiency is progress and that the frictionless delivery of information is an unambiguous good, provides no language with which to articulate it. It may well be that the cognitive uncanny valley is not a symptom of nostalgia. It is a perceptual signal. In evolutionary terms, perceptual signals that register mismatch between what something appears to be and what it actually is exists for good reasons. The discomfort is not the problem. The discomfort is the data.

Three Forms of Productive Friction

The argument so far has established that cognitive friction is not an obstacle to be overcome but a fundamentally constituent part of cognition. Possibility Studies frames this systemically: possibility-related phenomena depend on more than individuals, requiring a

wide ecosystem of human and non-human actors and their entanglement (Glăveanu, 2023; Harris, 2023; Ross, 2023a). Agency, from this perspective, is co-agency: the discovery and exploration of the possible depends on the relationship between person and environment rather than either of those in isolation. There are three distinct domains of such friction that human cognitive life has historically depended upon, and that FDEs are systematically redesigning away.

The first is the friction of materials. The physical properties of the tools and media through which we think have always shaped how we think. The resistance of clay or the grain of paper are not incidental textures or noise but cognitive partners, active participants in the assembly of understanding. Thinking happens not just in the brain but in the transactional space between brain, body, and material world (Vallée-Tourangeau, 2025). This is the space that Malafouris described when he demonstrates that the potter's knowledge of clay is not separable from the resistance clay offers the hand (Malafouris, 2008). As FDEs replace material engagement with linguistic and digital surfaces that offer no resistance, no grain and no unpredictability, something changes in the cognitive life that is assembled through them. However, our understanding of cognition has been so long cast in the mould of this dematerialised world, that we have little understanding of what we are in the process of losing.

The second is the friction of the body. Cognition is not a brain process that happens to occur in a body. It is an embodied process in which brain, body, and environment are constitutively entwined (Chemero, 2011; Fuchs, 2017; Gallagher, 2023). Physical exertion, gesture, movement, and full-body engagement with the world have all been shown to deepen cognitive encoding and support the kind of understanding that sedentary, screen-mediated environments cannot replicate (Goldin-Meadow, 2017; Lakoff & Johnson, 2010). As knowledge work migrates into FDEs, the body becomes a passenger rather than a participant in cognition and so the particular kind of knowing that is grown through bodily effort, through the coupling of physical and cognitive demand that Rosati (2017) traces back to the evolutionary roots of human intelligence, becomes progressively less available. The consequences of this are not yet fully understood, partly because the research paradigms most often used to study cognition in FDEs are themselves shaped by the model of mind that this paper has been questioning and they are therefore structurally unable to detect what is lost when the body is removed from the loop. What the embodied cognition literature (e.g. Newen et al., 2018) suggests, however, is that this loss is not trivial. The knowing that is assembled through bodily engagement has a different texture to the knowing assembled through purely

symbolic, screen-mediated processing. To remove the body may make cognition more efficient but it may also change what cognition is.

The third is the friction of other people. Perhaps the most cognitively consequential friction of all is the resistance of other minds. Human cognition evolved in deeply social contexts in which knowledge was acquired through dialogue, debate, and the challenge of different perspectives (Glăveanu, 2020; Tomasello, 1999). Possibility Studies argues that the possible is grounded in difference: minds defined by diversity and dialogue are open to new possibilities in ways that monological and fixated mindsets are not, and differences between people are a necessary condition for engaging the possible. The *unpredictability* of other people, their resistance to our ideas, their non-verbal signals of doubt, their refusal to understand until we explain better, constitutes a form of cognitive friction that structures and tests thinking in ways that no individualised system can replicate and one to which we are attuned. AI assistants are, by design, agreeable. The absence of interpersonal friction, like the absence of material and bodily friction, is something we feel before we can say what it is.

Conclusion

The goal is not to reject FDEs, which is as futile as it is misguided. It is to develop a clear understanding of what is being lost through this cognitive shift by taking the environment of thought seriously and understanding that human agents are always coupled with the world. Considering the importance of cognitive friction in this way can support design for FDEs which embrace friction. This is already happening; Microsoft has a Tools for Thought project which is drawing on interdisciplinary research to understand how AI can be designed to support rather than replace human thinking (Sarkar, 2025). The findings that better thinking happens when resistance and friction is built in creates artificially the friction that exists in traditional offloading onto material, bodes or people. The notion of an embodied AI is growing in education circles (Videla et al., 2025) and the pages of this journal has showcased the work of Ronald Beghetto on human and AI collaboration (Beghetto, 2023)

This is also not a call to design for greater cognitive efficiency. Research that points to the lack of efficiency in cognitive processes when enacted in FDEs misses the point. To frame the coupling of mind and FDEs as a story of damage or loss is already to concede too much to an unhelpful model of the human mind which relies on optimisation and indeed, the very model which drives this acceleration. Rather, I suggest there is no self-sufficient and ideal brain-based cognitive system which is either weakened or strengthened by technology because the metaphors of strengthening and weakening subscribe to a teleological fallacy of

cognitive progress. Instead, cognition is ongoing and has always been constituted through tools, materials, and social and physical environments. The mind is not an entity which the world acts upon, but a series of softly formed, transient, situationally assembled systems that perpetually connect, disconnect, and reconnect across brain, body, and world. The question is therefore not whether technology shapes cognition but the extent to which this particular reconfiguration is compatible with the embodied, material, and social substrates through which human cognitive life has historically been assembled and sustained. The efficient cognition that is being enacted is both unsatisfying and counter to the ongoing process of thinking in the world. To develop a narrative that runs counter to the narrative of progress and calls instead for the deliberate preservation of inefficiency, not-knowing, impasse, and discomfort through which understanding has always been assembled. It is a call to defend the right to be cognitively inefficient. This means taking the long way round, sitting with not-knowing, and following the detour that turns out to matter. These are not failures of cognitive performance. They are the conditions of it.

Possibility Studies insists that an ethics of possibility is imperative: we need sustained reflection on the scope and limits of our engagement with the possible and the consequences. The possible is also political. The construction of norms around who has the right to discover new possibilities and who is denied that right is a political act, and power relations are a key determinant in the dynamic of the possible. It is the case that proposals like the one in this editorial may run the risk of romanticising friction just as others may romanticise possibility. Forms of social friction can fall on those groups that are least able to bear it and AI promises to remove some forms of artificial gatekeeping that have long excluded different groups. However, it could also be the case that as efficiency becomes ever more possible, those with least access to the material, bodily, and social friction described in this paper will bear the greatest *cost* of a world redesigned for cognitive smoothness. The discomfort that increasing smoothness produces is not equally distributed, but neither will be its consequences. Taking that discomfort seriously as evidence is therefore not only a scientific responsibility. It is an ethical and political one.

References

Anderson, M. M., Crow,Nicole K., Mayberry, Ted, & Mitchell, Derrick. (2025). *AI can transform the classroom just like the calculator.* https://www.scientificamerican.com/article/ai-can-transform-the-classroom-just-like-the-calculator/

Arfini, S. (2019). *Ignorant cognition: A philosophical investigation of the cognitive features of not-knowing* (Vol. 46). Springer International Publishing.

Arfini, S. (2023). Serendipity and ignorance studies. In S. M. Copeland, W. Ross, & M. Sand (Eds), *Serendipity Science* (pp. 125–143). Springer International Publishing. https://doi.org/10.1007/978-3-031-33529-7_7

Bawden, D., & Robinson, L. (2009). The dark side of information: Overload, anxiety and other paradoxes and pathologies. *Journal of Information Science*, *35*(2), 180–191. https://doi.org/10.1177/0165551508095781

Bedard, J., Kropp, M., Hsu, M., Karaman, O.T., Hawes, J., & Rosen Kellerman, G. (2026, March 5). When using AI leads to 'brain fry'. *Harvard Business Review*. https://hbr.org/2026/03/when-using-ai-leads-to-brain-fry

Beghetto, R. A. (2023). A new horizon for possibility thinking: A conceptual case study of Human × AI collaboration. *Possibility Studies & Society*, *1*(3), 324–341. https://doi.org/10.1177/27538699231160136

Bellisari, A. (2008). Evolutionary origins of obesity. *Obesity Reviews*, *9*(2), 165–180. https://doi.org/10.1111/j.1467-789X.2007.00392.x

Chemero, A. (2011). *Radical embodied cognitive science*. MIT Press.

Chirayath, G., Premamalini, K., & Joseph, J. (2025). Cognitive offloading or cognitive overload? How AI alters the mental architecture of coping. *Frontiers in Psychology*, *16*, 1699320. https://doi.org/10.3389/fpsyg.2025.1699320

Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. Oxford University Press.

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, *58*(1), 7–19. https://doi.org/doi:10.1093/analys/58.1.7

Copeland, S. M. (2019). On serendipity in science: Discovery at the intersection of chance and wisdom. *Synthese*, *196*, 2385–2406. https://doi.org/10.1007/s11229-017-1544-3

De Jaegher, H. (2019). Loving and knowing: Reflections for an engaged epistemology. *Phenomenology and the Cognitive Sciences*. https://doi.org/10.1007/s11097-019-09634-5

de Saint Laurent, C. (2024). The possibility paradox: How technology both expands and contracts possibilities. *Possibility Studies & Society*, *2*(1), 3–8. https://doi.org/10.1177/27538699241232742

Erdelez, S., & Makri, S. (2020). Information encountering re-encountered: A conceptual re-examination of serendipity in the context of information acquisition. *Journal of Documentation*, *76*(3), 731–751. https://doi.org/10.1108/JD-08-2019-0151

Fuchs, T. (2017). *Ecology of the Brain: The phenomenology and biology of the embodied mind* (1st edn). Oxford University Press. https://doi.org/10.1093/med/9780199646883.001.0001

Gallagher, S. (2023). *Embodied and enactive approaches to cognition* (1st edn). Cambridge University Press. https://doi.org/10.1017/9781009209793

Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, *15*(1), 6. https://doi.org/10.3390/soc15010006

Glăveanu, V. P. (2020). A sociocultural theory of creativity: Bridging the social, the material, and the psychological. *Review of General Psychology*, 108926802096176. https://doi.org/10.1177/1089268020961763

Glăveanu, V. P. (2023). Possibility studies: A manifesto. *Possibility Studies & Society*, *1*(1–2), 275386992211275. https://doi.org/10.1177/27538699221127580

Gluckman, P. D. (with Hanson, M. A.). (2006). *Mismatch: Why our world no longer fits our bodies*. Oxford University Press.

Goldin-Meadow, S. (2017). Using our hands to change our minds. *Wiley Interdisciplinary Reviews: Cognitive Science*, *8*(1–2), e1368. https://doi.org/10.1002/wcs.1368

Gong, C., & Yang, Y. (2024). Google effects on memory: A meta-analytical review of the media effects of intensive Internet search behavior. *Frontiers in Public Health*, *12*, 1332030. https://doi.org/10.3389/fpubh.2024.1332030

Gruber, M. J., Gelman, B. D., & Ranganath, C. (2014). States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron*, *84*(2), 486–496. https://doi.org/10.1016/j.neuron.2014.08.060

Harris, D. X. (2023). Non-binary possibilities of creative agency. *Possibility Studies & Society*, *1*(1–2), 99–104. https://doi.org/10.1177/27538699231166485

Hołyst, J. A., Mayr, P., Thelwall, M., Frommholz, I., Havlin, S., Sela, A., Kenett, Y. N., Helic, D., Rehar, A., Maček, S. R., Kazienko, P., Kajdanowicz, T., Biecek, P., Szymanski, B. K., & Sienkiewicz, J. (2024). Protect our environment from information overload. *Nature Human Behaviour*, *8*(3), 402–403. https://doi.org/10.1038/s41562-024-01833-8

Kirsh, D. (1995). The intelligent use of space. *Artificial Intelligence*, *73*(1–2), 31–68. https://doi.org/10.1016/0004-3702(94)00017-U

Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). *Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task* (Version 2). arXiv. https://doi.org/10.48550/ARXIV.2506.08872

Lakoff, G., & Johnson, M. (2010). *Philosophy in the flesh: The embodied mind and its challenge to Western thought*. Basic Books.

Lee, H.-P. (Hank), Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & 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*, 1–22. https://doi.org/10.1145/3706598.3713778

Li, N. P., Van Vugt, M., & Colarelli, S. M. (2018). The evolutionary mismatch hypothesis: Implications for psychological science. *Current Directions in Psychological Science*, *27*(1), 38–44. https://doi.org/10.1177/0963721417731378

Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, *116*(1), 75–98. https://doi.org/10.1037/0033-2909.116.1.75

Malafouris, L. (2008). At the potter's wheel: An argument for material agency. In C. Knappett & L. Malafouris (Eds), *Material Agency* (pp. 19–36). Springer US. https://doi.org/10.1007/978-0-387-74711-8_2

Malafouris, L. (2013). *How things shape the mind: A theory of material engagement*. MIT Press.

Malafouris, L. (2015). Metaplasticity and the primacy of material engagement. *Time and Mind*, *8*(4), 351–371. https://doi.org/10.1080/1751696X.2015.1111564

Mori, M., MacDorman, K., & Kageki, N. (2012). The uncanny valley [from the field]. *IEEE Robotics & Automation Magazine*, *19*(2), 98–100. https://doi.org/10.1109/MRA.2012.2192811

Newen, A., Gallagher, S., & De Bruin, L. (2018). 4E cognition: Historical roots, key concepts, and central Issues. In A. Newen, S. Gallagher, & L. De Bruin, *The Oxford handbook of 4E cognition* (pp. 2–16). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780198735410.013.1

Pirolli, P., & Card, S. (1999). Information foraging. *Psychological Review*, *106*(4), 643–675. https://doi.org/10.1037/0033-295X.106.4.643

Pye, D. (1968). *The nature and art of workmanship* (Revised ed., repr). Herbert Press.

Rosati, A. G. (2017). Foraging cognition: Reviving the ecological intelligence hypothesis. *Trends in Cognitive Sciences*, *21*(9), 691–702. https://doi.org/10.1016/j.tics.2017.05.011

Ross, W. (2022). Heteroscalar serendipity and the importance of accidents. In W. Ross & S. Copeland (Eds), *The Art of Serendipity* (pp. 75–99). Palgrave MacMillan.

Ross, W. (2023a). Embracing difference and cultivating hope: The transformative potential of possibility studies. *Possibility Studies & Society*, *1*(1).

Ross, W. (2023b). The possibilities of disruption: Serendipity, accidents and impasse driven search. *Possibility Studies & Society*, 275386992311736. https://doi.org/10.1177/27538699231173625

Ross, W. (2024). Accidental thinking: A model of serendipity's cognitive processes. *Review of General Psychology*, 10892680241254759. https://doi.org/10.1177/10892680241254759

Ross, W., Glaveanu, V. P., & Chemero, A. (2024). The illusion of freedom: Towards an embedded notion of constraints. In C. Tromp, R. J. Sternberg, & D. Ambrose (Eds), *Constraints in Creativity* (pp. 166–181). BRILL. https://doi.org/10.1163/9789004707238

Ross, W., & Groves, M. (2023). Let's just see what happens: A qualitative case study of risk and uncertainty in the creative process. *The Journal of Creative Behavior*, *57*(2), 305–318. https://doi.org/10.1002/jocb.578

Sarkar, A. (2025, September). How to stop AI from killing your critical thinking. *TEDAI Vienna*. https://www.microsoft.com/en-us/research/publication/tedai-2025-artificial-intelligence-as-a-tool-for-thought/

Seifert, C. M. (2024). How impasse leads to insight: The prepared mind perspective. In C. Salvi, J. Wiley, & S. M. Smith (Eds), *The Emergence of Insight* (1st edn, pp. 84–112). Cambridge University Press. https://doi.org/10.1017/9781009244244.005

Seifert, C. M., Meyer, D. E., Davidson, N., Patalano, A. L., & Yaniv, I. (1994). Demystification of cognitive insight: Opportunistic assimilation and the prepared-mind hypothesis. In R. J. Sternberg & J. Davidson (Eds), *The nature of insight* (pp. 65–124). MIT Press.

Shalu, Verma, N., Dev, K., Bhardwaj, A. B., & Kumar, K. (2026). The cognitive cost of AI: How AI anxiety and attitudes influence decision fatigue in daily technology use. *Annals of Neurosciences*, *33*(1), 73–84. https://doi.org/10.1177/09727531251359872

Stadler, M., Bannert, M., & Sailer, M. (2024). Cognitive ease at a cost: LLMs reduce mental effort but compromise depth in student scientific inquiry. *Computers in Human Behavior*, *160*, 108386. https://doi.org/10.1016/j.chb.2024.108386

Sterelny, K. (2012). *The evolved apprentice: How evolution made humans unique*. The MIT press.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, *12*(2), 257–285. https://doi.org/10.1207/s15516709cog1202_4

Tomasello, M. (1999). *The cultural origins of human cognition.* Harvard University Press.

Vallée-Tourangeau, F. (2013). Interactivity, efficiency, and individual differences in mental arithmetic. *Experimental Psychology*, *60*, 302–311. https://doi.org/10.1027/1618-3169/a000200

Vallée-Tourangeau, F. (2025). Systemic creative problem-solving: On the poverty of ideas and the generative power of prototyping. *Systems*, *13*(3), 150. https://doi.org/10.3390/systems13030150

Videla, R., Penny, S., & Ross, W. (2025). 'If you can't dance your program, you can't write it': Challenges and Implications for AI in Education. *ACM Transactions on Computing Education*, 3759257. https://doi.org/10.1145/3759257

Voinea, C., Mann, S. P., Savulescu, J., & Earp, B. D. (2026). The calculator analogy: Epistemic virtues for using LLMs. *Technology in Society*, *85*, 103198. https://doi.org/10.1016/j.techsoc.2025.103198

Wilson, R. A. (2001). *Boundaries of the mind: The individual in the fragile sciences*. Cambridge University Press.

Yaqub, O. (2018). Serendipity: Towards a taxonomy and a theory. *Research Policy*, *47*, 169–179. https://doi.org/10.1016/j.respol.2017.10.007