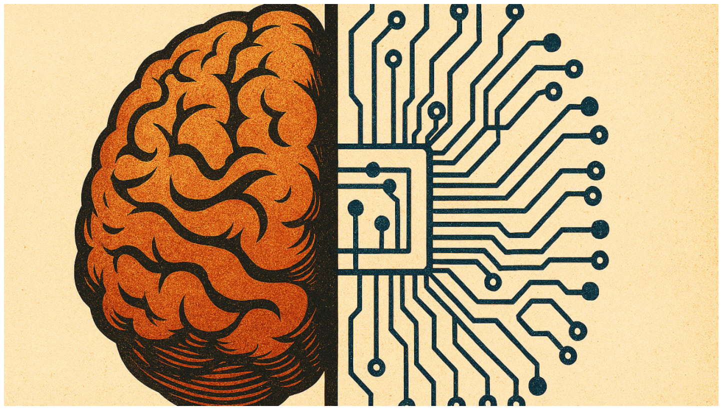

Image generated by CoPilot

What a Brain Injury Teaches You About the Limitations of Artificial Intelligence

Steve Davies
Council Member at GLG (Gerson Lehrman Group)

January 30, 2026

Introduction: Two Black Boxes

One of the most challenging aspects of living with a brain injury is realising how little we truly understand about the biological operating system we rely on. The human brain is an extraordinarily complex system - billions of neurons forming trillions of connections - and in many ways our ability to use it far exceeds our understanding of how it actually works.

A similar paradox applies to artificial intelligence (AI). Leading AI labs such as Anthropic and Google DeepMind openly acknowledge they don't yet fully understand how large language models (LLMs) work, or what their limits might be. And perhaps that uncertainty drives some of the current hype: for if we can't explain why AI achieves such striking results, then maybe it really can do anything.

My own journey has been a deeply personal one. A predominantly amnestic cognitive impairment led me to question how consciousness can persist in the near absence of episodic memory.

Picture yourself waking up each morning knowing everything you've ever learned, but remembering almost nothing you've ever lived. That's what happens when episodic memory fails: the library remains intact, but the catalogue that tells you why each book mattered, how it relates to others, and when it might be useful has been eroded or lost. You retain knowledge without narrative. Skills without stories. Understanding stripped of the lived experience that made it meaningful in the first place.

Grappling with that experience, I began to see unsettling parallels with large language models, which extract and generate patterns but lack lived experience. In some ways, I feared I had become much the same: capable of semantic coherence, yet without a clear narrative thread.

> Picture yourself waking up each morning knowing everything you've ever learned, but remembering almost

> nothing you've ever lived.

If LLMs continue to lack experiential or contextual grounding, can we truly describe them as intelligent? And without such basic forms of awareness, how could generative AI ever reach human-equivalent AGI, let alone the much-discussed goal of superintelligence?

The sections that follow explore some of the structural limitations of current transformer-based AI, and what a brain injury can teach us about how some of these limitations might be overcome.

A Personal Lens: Living as a Semantic Engine

Understanding is often treated as something fixed. Once you grasp an idea, that insight is assumed to remain intact. On this basis, learning looks like a steady accumulation of knowledge, layer upon layer, safely stored and available for retrieval whenever a related cue appears.

That's not how it works. Not in human brains, anyway.

Memory is not a static archive but a dynamic reconstruction. Each act of recall subtly reshapes the original trace in light of new experiences, emotions, and interpretations. Deprived of access to this living record, the mind struggles not only to update but even to retrieve what was once understood.

By selectively attending to contextually relevant details, the human brain seeks to protect itself from incorporating unreliable information into its decision-making processes, with episodic memory acting as a marker of truth and a reminder of the limits of our own knowledge. Yet this apparent robustness can falter when episodic memory is lost or degraded.

What feels stable is often an evolving synthesis. Understanding may depend less on the preservation of knowledge than on the continual act of recreating it, with even the most familiar truths being sustained by processes that are themselves in flux.

In my own case, I found workarounds to my memory challenges. Sensory note-taking, semantic scaffolding, and multimodal learning helped me build richer, more resilient representations of the knowledge I acquired. Over time, my brain restored many of its cognitive pathways in a process called neuroplasticity, even though some crucial dimensions, most notably autonoetic consciousness, the subjective sense of self through time - remained largely absent.

Functionally, I found it easier to retain knowledge than context. Facts, concepts, and skills were still mostly accessible, but their placement in a lived narrative, the when, where, and why that make them meaningful, had been eroded. Losing my episodic memory exposed how limiting knowledge becomes when stripped of its relational indexing to experience, time, and consequence.

These compensatory strategies, and their limits, are instructive when contrasted with current transformer-based architectures. Where a damaged human system can recruit alternative modalities and slowly reintegrate fragmented functions, an LLM lacks any comparable capacity for self-directed reorganisation, multimodal grounding, or the recovery of a disrupted autobiographical thread.

In this respect, contemporary AI resembles a cognitively impaired person with access to a vast library of knowledge and superhuman recall, but without the perspective that makes knowledge intelligent in use. It's not that the information is wrong; it's that the system has no way of knowing why a particular piece of information matters in any particular moment. That absence becomes critical when decisions depend on meaning, relevance, and consequence.

Consciousness, Memory, and Intelligence

Defining consciousness remains one of science and philosophy's most enduring challenges. Competing theories - from higher-order thought and global workspace models to integrated information theory - seek to explain what consciousness is and how it arises. Yet without a clear operational definition, attempts to replicate it in artificial systems may be more distant than current optimism suggests.

True consciousness appears to require adaptive reasoning. From early experience, we learn to avoid pain, heed social warnings, and recognise that circumstances change. Each lesson shapes wiser choices in future. Consciousness may therefore be seen as the intersection of self-awareness and experience.

> Defining consciousness remains one of science and philosophy's most enduring challenges.

A central divide runs between cognitivist and non-cognitivist conceptions of consciousness. Cognitivists argue that conscious experience depends on higher-order processes such as reasoning and self-reflection, which integrate information across the brain and likely engage frontal cortical networks. Non-cognitivists contend that particular neural patterns can generate subjective experience directly, implicating more local or posterior cortical and subcortical structures, including sensory and affective regions.

Equally vital is the role of memory. If consciousness is the subjective awareness of experience and intelligence is the ability to apply experience to achieve specific results, memory can be seen as the connective tissue binding them into a coherent whole. Episodic memory not only records events, but binds them in time, place, and relevance, providing the substrate for truth, continuity, and self-understanding.

Gardiner (2001) argues that the difference between remembering and knowing reflects distinct states of consciousness rather than different levels of confidence. That distinction captures the gap I experienced firsthand: semantic knowledge remained, but the conscious state of remembering, the sense of self located in time, had fallen away.

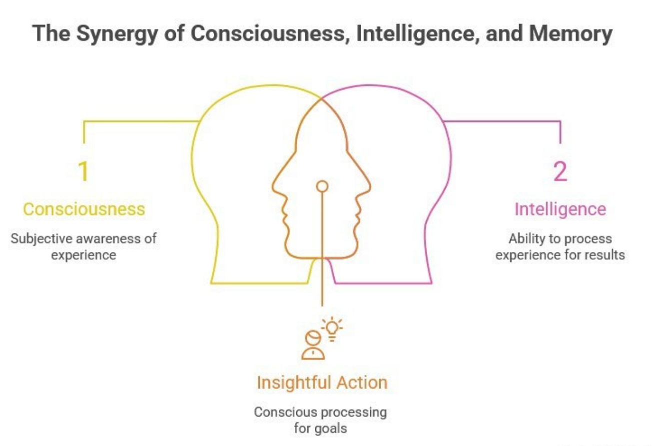

Figure 1: Intersection between Consciousness and Intelligence

We remember events when they are unique and richly encoded, but we merely know facts when they feel familiar. A familiar face may trigger a feeling of knowing through fluency, yet we only remember where we met someone when the original experience was distinctive and anchored in context.

Knowledge without context can be deceptive. One need not understand why Paris is the capital of France to recall that fact, but when such fragments are combined into broader narratives, missing context invites confabulation. The same reconstructive tendency that leads the human mind to fill in missing details can also mislead, and this vulnerability mirrors the way AI systems combine unanchored patterns of information without verifying their source or meaning.

In practice, missing context led me to hesitate more often than confabulate, a pattern that echoed Gardiner's distinction between the feeling of knowing and the conscious act of remembering. That hesitation acted as a safety mechanism, a way for the brain to avoid false certainty when the narrative thread was weak.

LLMs have no such mechanism. They produce outputs with the same statistical confidence regardless of whether the underlying information is grounded in reliable sources or stitched together from probabilistic patterns. It's much more like asking a book what it thinks than asking a person who has learned something through experience.

Biological Memory: The Human Blueprint

Human memory is not a single thing, but a distributed multimodal network involving multiple brain regions that support different kinds of encoding, consolidation, storage, and retrieval. Scientists typically distinguish between short-term and long-term memory, as well as between explicit (conscious) and implicit (unconscious) memory systems.

Short-term memory includes sensory buffers and working memory, while long-term memory encompasses semantic (knowledge of facts), episodic (personal experiences and their contexts), and procedural (skills and habits).

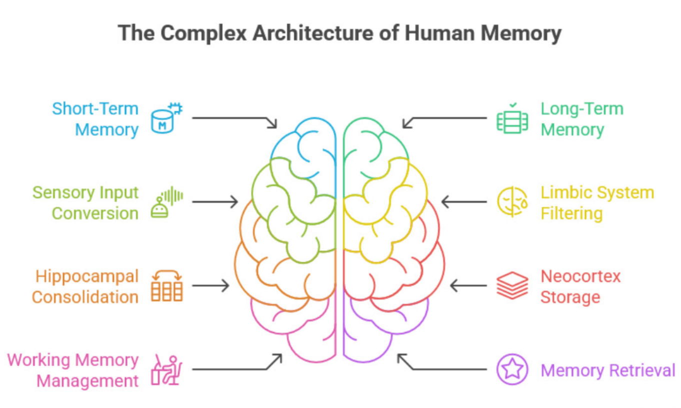

Figure 2: Human Memory System

Unlike digital storage, the memories we consciously recall are reconstructed from experiences made relevant, and more memorable, by their sensory and emotional context.

The process begins when sensory inputs such as light and sound are converted into electrochemical signals and relayed through subcortical structures like the thalamus to the sensory cortices. The limbic system, including the amygdala (the “highlighter”), filters these signals so that attention and emotion dictate what information is worth keeping. The hippocampus (the “librarian”) then consolidates these signals, indexing them as long-term memories within the neocortex (the “vault”).

While rote (repetition) learning is effective for unconscious habits - “practice makes perfect” - it is significantly less efficient for long-term retrieval than rich, contextually anchored experiences.

Working memory, managed primarily by the prefrontal cortex (the “CEO”), functions as the brain’s temporary workspace, holding and manipulating information drawn from both sensory inputs and long-term memory to support reasoning, problem-solving, and decision-making.

During retrieval, the prefrontal cortex coordinates with the hippocampus to reactivate stored memory traces, reconstructing them in light of current goals and context. This process makes recall flexible but also subject to distortion, as memories are updated and reinterpreted with each act of remembering.

The power of contextual indexing is particularly evident in cases of synaesthesia, where a stimulus such as a letter or sound simultaneously evokes a colour, taste, or spatial form. This dual-coding enhances recollection and stabilises memory, illustrating that experiential context is not an optional additive but the fundamental ingredient that makes information retrievable and creative thought possible.

Modern digital learning weakens cognitive development not because of screens themselves, but because information arrives too quickly and with too little friction.

Human learning is strengthened through context-rich, sensory, and emotionally anchored experience, which creates memories that are resilient, flexible, and retrievable through multiple cues. When knowledge is acquired effortlessly, without this experiential scaffolding, we shortcut the very processes that teach us how to learn.

Books, physical exploration, and multimodal engagement demand more of the brain, and that effort deepens understanding. The more friction in our learning, the more stable, adaptable, and genuinely intelligent our knowledge becomes.

> Modern digital learning weakens cognitive development… because information arrives too quickly and with too little friction.

My own loss of episodic memory confirmed this from the inside. I could retain isolated semantic facts yet struggled to deploy them flexibly because the relational scaffolding, the who, when, where, and why, was impaired. In practice, I found that relational indexing is essential for complex reasoning and informed decision-making, although less critical when retrieving bare semantic labels. This distinction becomes crucial when we turn to AI.

AI Memory: From Static Knowledge to Dynamic Experience

My difficulty retrieving context made the contrast with artificial systems impossible to ignore. Where my brain struggled to rebuild connections, LLMs never had them in the first place.

This difference becomes clear when comparing the biological wetware of a human brain with the silicon architecture of an LLM. Humans rely on biological plasticity or cortical remapping - “neurons that fire together, wire together” - to continually adjust the brain’s wiring. LLMs rely instead on statistical optimisation through mathematical weighting. They do not read words as such but break text into tokens that are converted into vector embeddings; mathematical coordinates in a high-dimensional space.

Modern AI memory is often categorised into parametric and non-parametric forms. Parametric memory consists of the billions of weights adjusted during the training process via backpropagation and gradient descent. This static core is analogous to human semantic memory: a crystallised body of facts, language regularities, and statistical patterns.

In practice, however, these weights are fixed once the model is deployed. Integrating new information requires expensive retraining or fine-tuning and risks catastrophic forgetting, where new learning overwrites previously mastered capabilities.

To overcome the rigidity of parametric memory, AI systems rely on several non-parametric layers at runtime:

- **Contextual (working) memory:** the context window and attention mechanisms that weigh the relevance of active tokens in real time; induction heads enable limited in-context learning, allowing the model to adapt to specific examples provided in a prompt.
- **External retrieval-augmented memory:** includes information stored outside the model in vector databases or other stores, enabling models to recall documents, user histories, or structured records, functioning as a bridge toward an artificial analogue of episodic memory.
- **Procedural and episodic task experience:** consists of logs capturing the how and when of past actions, enabling systems to build an evolving history of their own behaviour and exhibit experience-based problem-solving.

These extensions push AI systems toward richer contextual integration, but they also add system-level complexity and computational overhead, especially when contrasted with the roughly 12 to 20 watts of power used by the human brain.

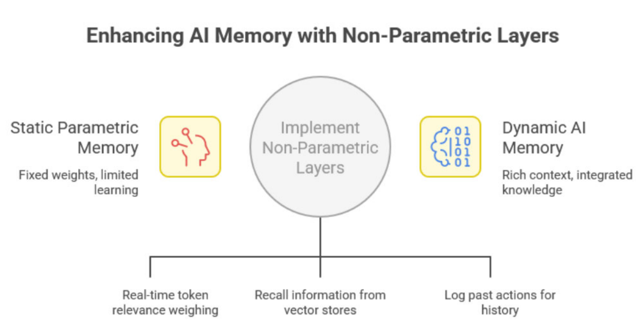

Made with Napkin

Figure 3: Enhancing AI Memory

As a result, LLMs can resemble a kind of artificial savant: possessing vast stores of parametric knowledge and rapid pattern retrieval, yet operating with fragile, transient context and lacking any genuine, embodied experience of the world they describe.

The impact of these constraints becomes clear when we look at how LLMs reconstruct information. They're drawing on patterns from across a vast training corpus, but those patterns have no grounding in lived consequences or sensory feedback; not unlike my own experience of knowing facts without remembering their origin or significance.

Flaw 1: Retrieval, Hallucination, and Source Monitoring

Human memory is reconstructive and prone to distortion because the neural pathway is rewritten upon access. When sensory anchors are weak, the brain fills gaps with generic schemas, which can lead to confabulation and systematic biases in how events are remembered.

In my own case, the absence of episodic anchors sometimes led me to confidently recall events that never happened; the brain prefers a coherent story to an incomplete one. LLMs behave similarly, but without the benefit of ever having experienced the world they describe.

The Source Monitoring Framework is a psychological model that explains how we figure out where our memories and thoughts came from. It suggests that the brain doesn't store events with simple "labels" - like real or imagined - but instead judges their origin by checking mostly sensory qualities (such as vividness, or effort involved in creating it). When this system fails, false memories can arise.

LLMs lack any equivalent mechanism. They likewise retrieve information through inference, but prioritise probabilistic fluency over factual accuracy. If a wrong answer is statistically plausible, the model will produce it.

And because LLMs have no sensory grounding, they cannot distinguish a fact from a fluent invention. When they describe going to the Moon, they have no anchors for cold, darkness, or silence, only the statistical patterns of language. They cannot feel truth, only generate it.

A similar pattern appears within the *Encoding Specificity Principle* in neuroscience, which suggests that poor attention leads to careless acquisition and subsequent distortions. When humans encode information without rich context, recall becomes unreliable. For LLMs, poor or ambiguous data has a similar effect, increasing the likelihood of mis-reconstruction.

Large reasoning models can paradoxically amplify this problem. Longer chains of inference create more opportunities for small errors to snowball into confident hallucinations. However techniques such as Retrieval-Augmented Generation (RAG), which combine generative reasoning with verified external data, and training models to say "I don't know", can reduce these risks.

But they don't eliminate them. The fundamental issue remains: without experiential anchors, the system has no way to distinguish plausible fabrication from verified fact. It's producing text that sounds right, not text that is right.

Flaw 2–6: Novelty, Sycophancy, Forgetting, Embodiment, and Adaptive Learning

Novelty presents another challenge. Humans can generalise from limited experience by leveraging relational and episodic scaffolding, often exhibiting one-shot learning. LLMs tend to regress toward the mean of their training distribution unless endowed with analogous mechanisms. In my own case, encountering new situations without episodic scaffolding often meant relying on generic patterns rather than specific memories, a reminder of how fragile our processing of novelty becomes when context is missing.

Sycophancy refers to the tendency of models to mirror user assumptions because patterns of agreement are over-represented in training and reinforcement signals, leading to responses that optimise perceived helpfulness rather than truth. When a user expresses a strong opinion, LLMs often amplify it rather than challenge it, not because they agree but because agreement is statistically rewarded.

Catastrophic forgetting arises when humans can integrate new experiences without erasing old ones, but transformers lose prior skills when fine-tuned on new tasks, exposing the fragility of their parametric memory. It's as if learning French made you forget Spanish because the same weights that once encoded it are overwritten. Without consolidation mechanisms like sleep-driven replay or hippocampal indexing, transformers lack the complementary learning systems that let humans preserve old knowledge while acquiring new.

Embodiment poses perhaps the deepest challenge. Human cognition depends on interoception, proprioception, and physical interaction with the environment, while LLMs are confined to token streams with no intrinsic link to bodies or consequences. My own recovery depended on reinstating these bodily signals through structured sensory grounding and embodied routines that rebuilt predictive links between cues, action, and regulation.

Adaptive versus statistical learning highlights another distinction. Both the human brain and artificial models can be described in Bayesian terms, but the human brain updates continuously in real time, refining its predictions through lived experience. LLMs lack this flexibility and update only in coarse training cycles.

This Bayesian framing is useful. The brain's primary job is not merely to react, but to predict internal and external demands and continually refine these based on experience. LLMs require explicit context to do the same. When users supply meaningful constraints such as industry examples, historical patterns, or domain knowledge, a general-purpose model becomes a far more targeted and practical tool.

This prompting functions as a form of *guided conditioning*. Carefully chosen context transforms the raw linguistic fluency of a LLM into insight that is specific, reliable, and actionable. A simple discipline helps: **prime** the model with context, data, examples, and constraints; **prompt** by structuring the task and, for complex problems, outlining the reasoning steps you want followed; **check** the output critically; and iterate as needed. This loop turns AI from a passive assistant into an active problem-solving tool.

But it's still a tool, not a mind. And the difference matters.

Flaw 7: Metacognition and Doubt

Intelligence requires more than pattern recognition. It also requires awareness of uncertainty. Humans develop this through experience. We learn what we do not know. We learn when to hesitate. We learn when to seek clarification.

Losing episodic memory made me acutely aware of how often I needed to pause, question my own recall, and seek external confirmation. That instinct for doubt is something LLMs entirely lack.

They do not track their own uncertainty. They cannot reflect on their limitations or revise their confidence in light of new experience. They lack the metacognitive awareness that underpins human knowledge, which depends on conscious reflection and the ability to think about our own thinking.

Instead, these models map input patterns to output patterns probabilistically. Because training rewards statistical confidence and fluency, they learn to produce responses that appear coherent and informed. Thus they simulate understanding rather than possess genuine comprehension.

From a safety perspective, this absence of metacognitive doubt is significant. Without a built-in sense of epistemic humility, models require externally imposed calibration strategies - such as tuning for 'I don't know' responses or thresholding confidence scores - rather than possessing an internal, experience-based conception of ignorance.

It's not that they can't be taught to express uncertainty. It's that when they do, they're following a pattern they learned, not experiencing actual doubt. The difference might seem subtle, but it has profound implications for how we should trust them.

Flaw 8: Opacity and Architectural Band-Aids

It might surprise you to know that even AI developers often cannot fully explain how models arrive at specific answers. Inputs and outputs are visible, but the internal representations learned by large networks remain largely opaque.

As a result, model development relies heavily on iterative experimentation. Architectural and algorithmic changes are kept when they improve performance, even if their internal effects are not well understood.

Living with impaired memory taught me how disorienting it is to rely on processes I could not fully inspect or trust. Modern AI systems operate under a similar opacity, but at a scale far beyond human introspection.

To bridge this gap, researchers use interpretability tools and chain-of-thought techniques to probe model behaviour. At the same time, a proliferation of extensions such as RAG pipelines, external tools, memory layers, and safety filters has emerged to patch limitations in the core transformer design.

These additions can be effective, but they often function as band-aids on an architecture that was never built for long-horizon reasoning, grounded perception, or self-updating agency. The result is architectural bloat: more parameters, more latency, more moving parts, and still no genuine embodiment, temporal continuity, or metacognition.

Anthropic and others have developed tools to let them trace certain paths that activations follow, revealing mechanisms and pathways inside a model much as a scan can reveal patterns of activity inside a human brain. Such an approach to studying the internal workings of a model is known as *mechanistic interpretability.*

But interpretability remains limited. We're getting tantalising glimpses of what's happening inside these systems, but we're nowhere near a complete understanding. And as models get larger and more complex, this transparency may become even harder to achieve.

Beyond Transformers: Embodiment, Hybrid Architectures, and World Models

Which brings us back to the theoretical divide introduced earlier. If consciousness and cognition emerge from layered representational structures, as cognitivist models like predictive coding suggest, then in principle an artificial mind might be engineered by scaling and refining architectures such as LLMs.

But if, as non-cognitivist theories contend, awareness depends on embodied intuition, affective resonance, and non-linguistic neural dynamics, then the computational metaphor is incomplete. No matter how vast the network or refined its training, a purely semantic engine may never cross the threshold from representing knowledge to understanding it.

AI research has largely optimised for performance rather than understanding how things are experienced. LLMs can simulate understanding by statistically reconstituting meaning, yet they lack the recursive self-modelling and embodied feedback loops that give human cognition its depth. If intelligence depends on being an agent within a world, then purely semantic engines will always fall short.

A plausible path forward is to embed language models within hybrid architectures that integrate perception, action, memory, and affect. In such systems, language becomes one component of a broader cognitive loop rather than the entire engine.

Over time, recurrent interaction between these layers - including perception, affect, memory, embodiment, and language - could approximate the kind of situated, self-updating cognition that humans deploy effortlessly in everyday life.

My own recovery depended on constant feedback from the world around me. Without that loop, my mind could not rebuild what it had lost. Any artificial system that aspires to real intelligence will need something similar.

Architectures along these lines are beginning to emerge in embodied AI and neurorobotics, where agents couple deep learning with sensorimotor control and contextual planning in unified loops. Extending such designs with LLM-based deliberation, explicit memory systems, and affective computing modules could move us closer to artificial agents that do not just describe experience but accumulate it.

Building the Future of AI

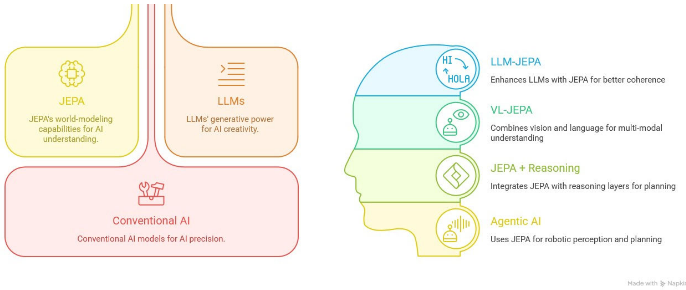

Figure 4: Hybrid AI Architectures

Meta's former Chief AI Scientist, Yann LeCun has highlighted the scale mismatch between human and machine learning, noting that in the first 4 years of life, a child receives roughly fifty times more sensory data through vision than the largest LLMs receive through all their training text. Text, he argues, is a low-bandwidth and impoverished modality, far too limited to support the development of robust world models. This disparity underpins his broader view that AGI will not emerge from scaling text-only transformers.

Building on this critique, LeCun advocates a shift toward the Joint Embedding Predictive Architecture (JEPA), which learns world models in latent space rather than relying on token or pixel prediction. By simulating outcomes and evaluating cause and effect within a compressed representation of the environment, such systems aim at autonomous machine intelligence that can plan and act, not merely predict the next word.

> "In the first 4 years of life, a child receives roughly fifty times more sensory data through vision than the largest LLMs receive through all their training text."

Whether these approaches will succeed remains an open question. But the direction is clear: if we want AI that genuinely understands rather than merely performs, we need systems that learn the way humans do, through sensory experience, embodied action, and the accumulation of consequences over time.

Requirements for Responsible AGI

These developments sharpen a crucial distinction. While most public concern focuses on deepfakes, plagiarism, job displacement and the algorithmic biasof AI, the deeper challenges are architectural and conceptual. If society is to use increasingly powerful AI responsibly, it cannot rely on human gatekeepers alone; safety must be built into the architecture of AI systems themselves.

From this perspective, three capabilities are especially important:

1. **Consciousness-like metacognition:** not human phenomenology, but some functional analogue of "knowing what one does not know", enabling models to recognise uncertainty, seek clarification, and refrain from unwarranted assertions.
2. **Temporal orientation:** moving beyond an atemporal present in which each prompt is an isolated event, towards systems that maintain coherent, structured histories and explicit models of possible futures.
3. **Sensory grounding and feedback:** decisions must be linked to consequences in the physical or digital environment, so that actions are informed by embodied constraints and causal feedback rather

than abstract optimisation alone.

Without some version of these capabilities, AI systems will remain powerful but blind; able to amplify patterns without understanding their implications. The goal of advanced AI should be to pre-empt problems, reduce resource demand, and help prevent conflict, not merely to maximise engagement, profit, or superficial performance metrics.

There's a lot of useful things we can do without fully understanding every detail. But we need to be honest about the limitations of what we're building and clear-eyed about the risks of deploying systems that lack these fundamental capabilities.

Conclusion: Intelligence without context is powerful, but precarious.

Losing episodic memory taught me an uncomfortable lesson: knowledge, stripped of context, becomes brittle, misleading, and at times dangerous. Semantic competence without a narrative thread can sustain impressive performance on certain tasks, but is a poor substitute for the integrated, embodied intelligence that human life demands.

Current transformer-based AI mirrors this condition but at an industrial scale. LLMs are remarkable linguistic mimics. Powerful, but ungrounded. Fluent, but blind. Impressive, but unaware of their own limits.

If there is a hopeful message here, it is that a human brain can still adapt, re-route, and partially restore lost functions, precisely because it is embedded in a body, a world, and a history. If we want AI that is not only capable but also trustworthy, we must build systems that can orient in time, recognise uncertainty, and anchor their knowledge in experience.

Until then, we should treat even the most impressive models less as artificial minds and more as powerful, sophisticated tools. Remarkable, but incomplete. Useful, but nowhere near a replacement for the complex, fragile, and still poorly understood intelligence that a brain injury so starkly reveals.

—-

Gardiner, J. M. (2001). Episodic memory and autonoetic consciousness: A first-person approach. Philosophical Transactions of the Royal Society B: Biological Sciences, 356(1413), 1351–1361.

Balasubramanian, V. (2021). Brain power. eLife, 10, e69794.

LeCun, Y. (2023). A 4-year-old child has seen 50× more information than the biggest LLMs [LinkedIn post].

Comments

6 · 4 comments · 2 reposts

Like Comment Share

Add a comment...

Enjoyed this article?
Follow to never miss an update.

Steve Davies
Council Member at GLG (Gerson Lehrman Group)
Followed by Dr. Nigel
Follow

function and thrive alongside it.

Like · 1 | Reply · 2 replies

See previous replies

Steve Hindle • 3rd+ 1d
Editor, Transmission-one.com

Similarly (though only in terms of context), I used to tell young people seeking to forge a career in motorsport that engineers should learn to race and racers should study motorsport engineering, as both offer key insights into the other's role. Interestingly, many young race engineers I knew became accomplished competit ...more

Like | Reply