The Biological Prerequisite for Artificial General Intelligence: Why Probabilistic Computation Cannot Produce Cognition

Murat Atilgan
School of Computing, Engineering and Physical Sciences
University of the West of Scotland

ma@caledonianai.co.uk

Abstract

This paper argues that Artificial General Intelligence (AGI) and its theoretical successor, Artificial Superintelligence (ASI), are fundamentally unachievable through probabilistic computation alone, regardless of model scale, architectural innovation, or computational investment. We establish that all current AI systems—including large language models, diffusion models, and reinforcement learning agents—operate through statistical pattern matching over structured data representations. While these systems produce outputs that superficially resemble cognitive behaviour, they lack the defining properties of cognition: embodied experience, temporal continuity, homeostatic self-regulation, and phenomenal consciousness. We argue that these properties are not emergent features of sufficient computational complexity but are intrinsic to biological neural substrates operating through electrochemical processes that cannot be replicated through digital simulation. The paper presents two logically exhaustive paths to AGI: (1) direct bidirectional integration between artificial systems and biological neural tissue, or (2) complete replication of human neural architecture at biological fidelity. Both paths require breakthroughs in neuroscience, bioengineering, and materials science—not in software or computational scaling. We conclude that the prevailing industry narrative of achieving AGI through larger models and more compute represents a category error of historic proportions, and that genuine progress toward AGI requires redirecting research investment toward neurotechnology and biological-artificial integration.

Keywords: Artificial General Intelligence, Neurotechnology, Brain-Computer Interfaces, Probabilistic Computation, Cognition, Biological Neural Integration, Large Language Models, Consciousness, Neuro-AI Convergence

1 Introduction

The artificial intelligence industry is currently dominated by a single narrative: that Artificial General Intelligence will emerge from scaling existing computational approaches. Larger models, larger datasets, more parameters, more compute—the prevailing assumption is that sufficient scale will produce a qualitative transition from narrow pattern matching to genuine cognition. Major laboratories have staked billions of dollars and institutional credibility on this thesis. Public discourse increasingly treats AGI as an imminent engineering milestone rather than an unsolved scientific problem.

This paper challenges that narrative at its foundation.

The argument proceeds from a straightforward logical chain. All current AI systems, without exception, operate through probabilistic inference over data representations. Probabilistic inference—regardless of its scale, sophistication, or architectural implementation—computes statistical relationships between input patterns and output distributions. This is what these systems do. It is all they do. And it is categorically insufficient to produce cognition.

Cognition, as exhibited by biological neural systems, involves properties that have no probabilistic analogue: embodied sensory experience, persistent identity across time, homeostatic self-regulation, emotional valence that shapes reasoning, and phenomenal consciousness. These are not emergent features that appear when a statistical model crosses some threshold of complexity. They are properties of biological substrates operating through electrochemical processes fundamentally different in kind from digital computation.

The implications are direct. AGI cannot be achieved through software alone. It cannot be achieved through computational scaling. It requires physical integration with—or faithful replication of—biological neural architecture. This is not a pessimistic position; it is a redirective one. It suggests that the path to AGI runs through neurotechnology, not through larger GPU clusters.

2 What Current AI Systems Actually Do

Before arguing what AI systems cannot achieve, it is necessary to establish precisely what they do achieve—and through what mechanism.

2.1 The Mathematical Reality

Every modern AI system operates through a common mathematical framework. A model receives an input x, transforms it through learned parameters θ, and produces an output distribution $P(y|x,\theta)$ from which a response is sampled or selected. In large language models, x is a sequence of tokens, θ comprises billions of weight matrices, and y is a probability distribution over the vocabulary for the next token. The training process adjusts θ to minimise a loss function—typically cross-entropy—over a corpus of examples.

This is, at its mathematical core, statistical regression at extraordinary scale. The model learns conditional probability distributions from data. It does not learn "meaning." It does not develop "understanding." It computes which output sequences are statistically most likely given input sequences, as determined by patterns in its training data.

2.2 The Sophistication Trap

The outputs produced by modern LLMs are remarkably sophisticated. They can generate coherent essays, solve mathematical problems, write functional code, and engage in what appears to be nuanced reasoning. This sophistication creates a powerful illusion: if the output resembles cognition, perhaps cognition is what produced it.

This inference is logically invalid. A sufficiently detailed map of London is not London. A sufficiently accurate weather model does not produce rain. A system that generates text statistically indistinguishable from text produced by a thinking being is not thereby a thinking being. The quality of the output tells us about the fidelity of the statistical model, not about the presence of cognition in the system producing it.

2.3 The Scaling Hypothesis and Its Limits

The dominant industry position holds that cognitive capabilities will emerge from sufficient scale—that there exists some threshold of parameters, data, or compute beyond which the system transitions from sophisticated pattern matching to genuine understanding. This hypothesis is unfalsifiable in practice (one can always argue that more scale is needed) and unsupported by any theoretical framework that would predict such a transition.

What scaling demonstrably produces is higher-fidelity interpolation within the training distribution. Models become better at the statistical task they were designed to perform. They do not become something other than statistical systems. A linear regression with a trillion parameters is still a linear regression. The mathematical operations do not change their fundamental character because they are performed at larger scale.

3 What Cognition Actually Requires

Having established what current AI systems do, we now examine what cognition—the target capability of AGI—actually entails. The gap between these two descriptions constitutes the central argument of this paper.

3.1 Embodied Experience

Biological cognition does not operate on abstract data representations. It operates on sensory experience—photons striking retinal cells, pressure waves vibrating cochlear hair cells, chemical molecules binding to olfactory receptors, mechanical deformation activating proprioceptive neurons. Each sensory modality produces not merely data but *experience*: the subjective quality of seeing red, hearing a minor chord, tasting salt.

An AI system processing the text string "the sunset was beautiful" performs statistical operations on token embeddings. A human experiencing a sunset undergoes photochemical cascades in retinal neurons, generates emotional responses through limbic system activation, modulates hormonal state through hypothalamic signalling, and forms episodic memories through hippocampal encoding. These are not two different descriptions of the same process. They are categorically different processes that happen to produce overlapping behavioural outputs.

No quantity of text processing can produce embodied experience, because text is a symbolic representation of experience, not experience itself. Training a model on every description of every sunset ever written does not give the model the experience of a sunset any more than reading a cookbook gives the reader the taste of a meal.

3.2 Temporal Continuity and Persistent Identity

A human being maintains continuous identity across time. Memories from childhood inform adult decisions. Personality develops through accumulated experience. The sense of self persists through sleep, anaesthesia, and even significant neurological injury.

An LLM has no temporal continuity whatsoever. Each inference call begins from the same initial state. The "memory" within a context window is not memory in any cognitive sense—it is a sequence of tokens that influence probability distributions for the current generation step. When the context window is cleared, nothing persists. There is no accumulated self, no developmental trajectory, no identity that evolves through experience.

Retrieval-augmented generation and persistent memory systems store and retrieve text—they do not create temporal continuity of experience. The distinction is between a being that remembers and a database that is queried.

3.3 Homeostatic Self-Regulation

Biological neural systems do not exist in isolation. They are embedded in bodies with metabolic needs, hormonal cycles, circadian rhythms, immune responses, and autonomic regulation. Cognition is not separable from this embodiment. Hunger affects decision-making. Fatigue degrades reasoning. Emotional states—mediated by neurotransmitter balances, hormonal levels, and autonomic nervous system activity—fundamentally shape what and how a biological system thinks.

An AI system has no homeostatic state. It does not get tired. It does not get hungry. It has no emotional valence that shifts its reasoning priorities. It processes each query with identical computational resources regardless of context. This is often presented as an advantage—consistency, reliability, tirelessness. But it is precisely this absence of embodied self-regulation that prevents AI systems from exhibiting the adaptive, context-sensitive, self-modifying behaviour that characterises general intelligence.

General intelligence is not the ability to solve any problem presented in text. It is the ability to determine which problems matter, how urgently they matter, and how to allocate finite cognitive resources under conditions of uncertainty, fatigue, emotional pressure, and competing biological imperatives. This capability is intrinsically biological.

3.4 Phenomenal Consciousness

The most fundamental gap between probabilistic computation and cognition is phenomenal consciousness—the subjective experience of "what it is like" to be a thinking entity. A human does not merely process information about pain; a human *experiences* pain. This experience has causal power: it motivates behaviour, shapes learning, and alters future decision-making in ways that pure information processing cannot replicate.

Whether AI systems can or will achieve consciousness is a deep philosophical question. This paper takes the more conservative and empirically grounded position: there is currently no evidence that probabilistic computation over digital substrates produces phenomenal consciousness, no theoretical framework that predicts it should, and no mechanism by which it could. The burden of proof lies with those claiming that sufficient statistical computation will produce subjective experience—a claim for which no supporting evidence exists.

4 The Category Error at the Heart of AGI Discourse

The prevailing AGI narrative commits a category error: it conflates the ability to produce cognitive-seeming outputs with the presence of cognition. This error has a precise structure.

4.1 Correlation Is Not Cognition

Current AI systems learn statistical correlations in data. They learn that certain token sequences follow other token sequences with certain probabilities. When these correlations are sufficiently comprehensive—when the model has learned enough of the statistical structure of human-generated text—the outputs become indistinguishable from those produced by cognitive beings.

But learning the statistical structure of cognitive outputs is not the same as performing cognition. A model that has learned the probability distribution over all human-generated text about quantum mechanics has not thereby learned quantum mechanics. It has learned which sequences of words about quantum mechanics are statistically likely. These are fundamentally different achievements.

4.2 The Chinese Room, Revisited

Searle's Chinese Room argument remains relevant despite decades of attempted refutation. A system that manipulates symbols according to rules—no matter how complex those rules, no matter how vast the rule set—does not thereby understand what those symbols mean. Understanding requires a semantic relationship between the symbol and its referent that cannot be captured by syntactic operations alone.

Modern LLMs are Chinese Rooms operating at trillion-parameter scale. The room is larger. The rule book is more comprehensive. The outputs are more convincing. But the fundamental character of the operation—syntactic manipulation without semantic understanding—has not changed.

4.3 The Turing Test Fallacy

If a system produces outputs behaviourally indistinguishable from those of a cognitive being, have we achieved AGI? The answer is no, for the same reason that a perfect flight simulator is not flight. Behavioural equivalence at the output level does not entail mechanistic equivalence at the process level. And it is the process—not the output—that defines cognition.

AGI is not defined as "a system that produces human-like text." It is defined as a system that possesses general intelligence—the capacity for understanding, reasoning, learning, and adapting across arbitrary domains. Producing text that *describes* understanding is not understanding. Generating text that *resembles* reasoning is not reasoning. The distinction is not subtle; it is categorical.

5 The Biological Bridge: Two Paths to AGI

If probabilistic computation alone cannot produce cognition, what can? This paper identifies two logically exhaustive paths, both requiring biological neural integration.

5.1 Path 1: Direct Bidirectional Neural Integration

The first path involves creating a physical interface between artificial computational systems and biological neural tissue. This is the direction pursued by brain-computer interface (BCI) companies including Neuralink, Synchron, Merge Labs, and Blackrock Neurotech.

Current BCI technology is overwhelmingly unidirectional: it reads neural signals and translates them into digital commands. This is valuable for medical applications—restoring motor control, enabling communication for locked-in patients—but it does not constitute the integration required for AGI.

The integration required would be *bidirectional at cognitive fidelity*: artificial systems would not merely read neural signals but participate in neural computation. Information would flow from biological neurons to silicon and back, with the artificial system contributing to the cognitive process rather than merely observing it.

This requires solving problems that do not currently have solutions:

- **Signal fidelity**: Current non-invasive BCIs capture aggregate electrical activity (EEG) at millimetre-to-centimetre spatial resolution. Cognitive-level integration would require single-neuron resolution across large cortical areas.
- **Bidirectional communication**: Writing information into neural tissue with sufficient precision to participate in cognitive processes, not merely stimulate gross motor responses.
- **Temporal synchronisation**: Biological neural computation operates on millisecond timescales with electrochemical signalling. Digital computation operates on nanosecond timescales with electronic signalling. Bridging this temporal gap without introducing latency that disrupts cognitive coherence is an unsolved engineering challenge.
- **Biocompatibility**: Long-term stable interfaces between silicon and neural tissue that do not degrade, provoke immune responses, or damage the biological substrate.

These are not software challenges. They are problems in neuroscience, materials science, bioengineering, and physics. No amount of algorithmic improvement addresses them.

5.2 Path 2: Complete Neural Architecture Replication

The second path involves replicating the human neural system at biological fidelity—not simulating it mathematically, but reproducing its physical structure and electrochemical dynamics in a substrate capable of supporting the same processes.

The scale of this challenge is extraordinary. The human brain contains approximately 86 billion neurons, each with an average of 7,000 synaptic connections, yielding roughly 600 trillion synapses. Each synapse operates through electrochemical processes involving dozens of neurotransmitter types, receptor subtypes, and intracellular signalling cascades. The state of each neuron is influenced by its chemical microenvironment, including hormonal concentrations, metabolic substrate availability, and local immune activity.

Replicating this system would require not merely modelling the connectivity graph—which is itself an unsolved mapping problem (the "connectome")—but reproducing the electrochemical dynamics at each node. A mathematical simulation of a neuron is not a neuron, for the same reason that a mathematical simulation of water is not wet. The computation must occur in a physical substrate that supports the relevant causal processes.

This path is not impossible in principle, but it requires advances in synthetic biology, neuromorphic engineering, and computational neuroscience that are decades away at minimum.

5.3 Why Software Cannot Substitute for Either Path

Both paths require physical breakthroughs—hardware, materials, biological engineering. Neither can be achieved through software innovation alone. This is the central claim of this paper, and it follows directly from the argument established in Sections 2–4.

If cognition requires embodied experience, then a system must have a body—or be connected to one. If cognition requires temporal continuity, then a system must have persistent physical state that evolves through experience. If cognition requires homeostatic self-regulation, then a system must be embedded in a regulatory framework with metabolic, hormonal, and autonomic dynamics. If cognition requires phenomenal consciousness, then a system must operate on a substrate that supports subjective experience.

Software runs on digital substrates that provide none of these properties. More sophisticated software does not change the substrate. A faster processor does not introduce embodiment. A larger model does not create experience. The limitation is not computational; it is physical.

6 The Role of Data: Necessary But Fundamentally Insufficient

This paper does not argue that data is unimportant. Structured, high-quality, domain-specific data is essential for every current and future AI system. The author's own work on domain-isolated LLM architectures demonstrates that data quality and structural discipline are the primary determinants of inference reliability within existing probabilistic frameworks.

But data—no matter how clean, how structured, how comprehensive—is a *representation* of the world, not the world itself. An AI system trained on the complete corpus of human knowledge about neuroscience does not thereby understand neuroscience. It has learned the statistical structure of text about neuroscience. The gap between these two achievements is the gap between correlation and cognition.

Data serves as the foundation for pattern matching at any scale. For AGI, data would need to be not merely *about* biological experience but *derived from* biological experience—real-time neural signals from actual nervous systems, integrated bidirectionally into computational processes. This is the bridge between structured data (which current systems handle) and embodied experience (which current systems cannot access).

The distinction is precise: current AI systems consume data *about* the world. A cognitive system would need to consume—and generate—data *from* the world, through sensory apparatus grounded in biological or biologically-equivalent substrates.

7 Implications for Research and Industry

7.1 The Misallocation of Resources

If the arguments presented in this paper are sound, the current allocation of AI research investment represents a historic misallocation. Billions of dollars are directed toward scaling probabilistic models—building larger GPU clusters, training larger transformers, curating larger datasets—in pursuit of AGI through a methodology that cannot, in principle, achieve it.

This is not to say that current AI research is valueless. Large language models, computer vision systems, and reinforcement learning agents are extraordinarily useful tools. They will continue to transform industries, accelerate scientific research, and augment human capability. But they will do so as *tools*—sophisticated, powerful tools operating through statistical inference—not as cognitive agents.

The research investment required to pursue AGI through biological integration—neurotechnology, brain-computer interfaces, synthetic biology, neuromorphic computing—is a fraction of what is currently spent on scaling existing architectures. Redirecting even a portion of current AI investment toward these pathways would accelerate genuine progress toward AGI while maintaining the practical benefits of continued narrow AI development.

7.2 The Neurotechnology Pathway

The emerging neurotechnology sector is the most plausible pathway toward the biological integration that AGI requires. Brain-computer interface companies are developing the physical infrastructure for neural-digital communication. Neuromorphic chip designers are building computational substrates that more closely mirror biological neural dynamics. Synthetic biology researchers are exploring the creation of biological neural tissue outside the human body.

These efforts are currently motivated primarily by medical applications—restoring motor function, treating neurological disorders, augmenting cognitive performance. But the infrastructure they build—high-bandwidth neural interfaces, bidirectional communication protocols, bio-digital signal processing pipelines—is precisely the infrastructure that a genuine AGI programme would require.

7.3 The Data Infrastructure Layer

When biological breakthroughs arrive, the neural data flowing through bio-digital interfaces will require processing infrastructure: domain isolation to prevent signal contamination between neural subsystems, clean data pipelines for real-time neural signal processing, retrieval architectures for grounding artificial computation in biological context, and quality frameworks ensuring signal fidelity across the interface boundary.

These requirements map directly onto the challenges already being addressed in current AI systems engineering—suggesting that expertise in domain-isolated architectures, clean data methodology, and contamination prevention will become foundational skills for neuro-AI integration.

8 Anticipated Objections

8.1 "Emergence Will Produce Cognition at Sufficient Scale"

This objection holds that cognitive properties will emerge spontaneously from sufficient computational complexity, as consciousness allegedly emerged from sufficient biological complexity. The response is twofold. First, biological complexity is not merely computational complexity—it includes electrochemical dynamics, embodiment, and homeostatic regulation that have no digital analogue. Second, the emergence claim is unfalsifiable: no proposed threshold of scale would definitively disprove it. Unfalsifiable claims are not scientific hypotheses; they are articles of faith.

8.2 "Functional Equivalence Is Sufficient"

This objection argues that if a system produces outputs functionally equivalent to those of a cognitive being, it is, for all practical purposes, cognitive. The response is that functional equivalence at the output level does not entail mechanistic equivalence at the process level. A calculator produces functionally equivalent mathematical outputs to a human mathematician; it is not thereby performing mathematics in the cognitive sense. AGI is defined by the nature of the process, not merely the quality of the output.

8.3 "Biological Chauvinism"

This objection accuses the position of assuming cognition requires biological substrates specifically, rather than any substrate with the right functional organisation. The response is precise: this paper does not claim that cognition requires carbon-based biology specifically. It claims that cognition requires a substrate capable of supporting embodied experience, temporal continuity, homeostatic regulation, and phenomenal consciousness. Currently, the only known substrate with these properties is biological neural tissue. If a synthetic substrate demonstrably supporting these properties were created, it would satisfy the requirements of this argument. But creating such a substrate is itself a biological engineering challenge, not a software engineering challenge.

8.4 "Current AI Systems Show Signs of Understanding"

This objection points to specific AI behaviours—in-context learning, chain-of-thought reasoning, apparent creativity—as evidence of nascent cognition. The response is that each of these behaviours has a complete explanation in terms of statistical pattern matching over training distributions. In-context learning is conditional probability estimation. Chain-of-thought reasoning is sequential token generation following patterns observed in training data. Apparent creativity is novel recombination of statistical patterns. Explaining these behaviours does not require positing cognition; therefore, by parsimony, cognition should not be posited.

9 Conclusion

The path to Artificial General Intelligence does not run through larger language models, larger datasets, or larger GPU clusters. It runs through the intersection of neuroscience and engineering—through physical systems capable of bridging biological neural computation and artificial information processing.

Current AI systems are, and will remain, probabilistic inference engines. They are powerful. They are useful. They are transforming how humans work, communicate, and discover. But they are not cognitive, and scaling them will not make them cognitive. The mathematical operations that define their function—matrix multiplication, gradient descent, softmax normalisation, attention computation—do not change their fundamental character at any scale. A trillion-parameter model performing conditional probability estimation is performing conditional probability estimation.

AGI requires either direct bidirectional integration with biological neural tissue or complete replication of neural architecture at biological fidelity. Both paths demand breakthroughs in neurotechnology, materials science, and bioengineering. Neither path can be shortcut through software alone.

The industry's current trajectory—scaling probabilistic computation toward an assumed cognitive threshold—represents a category error. It confuses the map for the territory, the simulation for the phenomenon, the statistical approximation of cognition for cognition itself. Recognising this error is not cause for pessimism; it is the prerequisite for directing resources toward approaches that can actually succeed.

The question is not whether artificial systems will achieve general intelligence. The question is whether we will pursue it through the only pathways that can deliver it.

> *Cognition is not computation. It is computation embedded in biology. Until artificial systems are embedded in—or faithfully replicate—biological neural substrates, they will remain what they are: extraordinarily sophisticated tools for statistical inference. Powerful, useful, transformative—but not thinking. Not yet.*

Author Note

The arguments presented in this paper were developed through the author's work on QAI Lab, a domain-isolated LLM orchestration platform. Extensive operational testing across 50+ model configurations revealed the fundamental limitations of probabilistic inference—limitations that motivated this investigation into what would actually be required to transcend them. The author's background in both computer science and international business operations informs the pragmatic, engineering-oriented perspective applied throughout.

Declaration

The author declares no conflicts of interest. This paper received no external funding. All arguments represent the author's independent analysis.