The Quantum Necessity Conjecture for Artificial General Intelligence

A position on the physical requirements of general intelligence

Andrew Koski

5 November 2026 | Working paper · Version 1.0

Abstract

I conjecture that artificial systems achieving human-level general intelligence within realistic physical resource limits will require a functional quantum component integrated with classical computation. The motivation is that intelligence may depend on physical processes whose computational role is lost when the brain is treated only as a network of classical information-processing units. That possibility does not follow merely from the fact that brains are biological or that digital computers use binary representations. A credible necessity claim must connect a specific physical mechanism to an indispensable cognitive capability and show why classical alternatives cannot supply that capability within stated constraints. This paper separates quantum advantage, practical quantum necessity, and impossibility of classical AGI in principle. It examines relevant work on neural dynamics, decoherence, nuclear-spin proposals, and collective optical effects in biological structures, then develops a framework for testing the conjecture. The framework requires independent evidence of biological relevance, comparisons against strong classical alternatives, and explicit conditions for rejecting the claim. The result is a research position and a proposed standard of evidence, not an experimental demonstration or an impossibility proof.

Keywords: artificial general intelligence; quantum computation; quantum cognition; computational neuroscience; resource constraints; falsifiability

1. The position

My position is that the physical basis of intelligence deserves a larger role in the AGI debate. We can keep improving algorithms and increasing computational resources while still asking whether something essential has been left out of the model. If a capability depends on a particular physical process, reproducing its outward description may be much easier than reproducing the process that makes it work.

I expect a successful account of general intelligence to connect computation with the physical mechanisms that make flexible learning possible. More specifically, I conjecture that a useful quantum resource will prove necessary for human-level general intelligence under realistic constraints. In an engineered system, that would mean quantum processing working with classical memory, learning, and control.

The intuition that a human brain is more than a collection of binary switches motivates this position. It does not establish it. A scientific argument has to explain which properties matter, what they do, and why another implementation cannot provide the same capability. Without those steps, the argument stops at an observation about hardware.

There are two routes to investigating the conjecture. Neuroscience can identify a quantum process that contributes to cognition. Computer science can identify an indispensable cognitive operation for which quantum computation provides an otherwise unavailable resource advantage. Either route could inform the other. Neither has to assume that an artificial mind must reproduce every detail of a human brain.

This is a position paper. It offers no new experimental data. Its purpose is to state a prediction clearly enough that future results can be judged against the original claim, including results that contradict it.

2. What is being claimed

Here, AGI means broad competence in learning and solving unfamiliar problems across substantially different domains, with transfer between them. A system should adapt when conditions change, retain useful knowledge, and carry out extended tasks without a new purpose-built solution for each one. This is a functional definition. Morris and colleagues similarly distinguish the breadth and depth of capabilities when describing levels of AGI; their framework helps make performance claims more specific. [1]

Consciousness is a separate question. A system might satisfy a functional intelligence criterion without settling whether it has subjective experience. If evidence eventually shows that consciousness is necessary for a particular cognitive capability, that link will need its own argument. Defining AGI as conscious only after a classical system passes an agreed test would change the claim being tested.

A classical system, in this discussion, processes information through classical states and operations; it may be deterministic, probabilistic, or analog. Its transistors can depend on quantum physics at the material level and still implement classical computation. A hybrid system deliberately uses nonclassical information processing, such as coherent quantum interference, as part of the computation. Simply adding a quantum device, or sourcing random numbers from one, does not establish that quantum computation is necessary for its intelligence.

Claim	What would have to be established
Quantum advantage	A quantum component improves a relevant task under a fair comparison.
Practical quantum necessity	The agreed capability is attainable within specified resources using a quantum component, while classical alternatives cannot attain it within those resources.
Impossibility in principle	No classical system can realize the defined intelligence, even with arbitrarily large finite resources.

The central conjecture concerns practical necessity. The stronger claim that classical AGI is impossible in principle carries a much larger burden of proof. A useful quantum speedup would not, by itself, establish either necessity claim.

3. Why the binary argument is insufficient

A brain is an electrochemical system. Neural activity involves changing voltages, ion flows, and dynamics over time. Hodgkin and Huxley showed how membrane currents and nerve excitation could be described mathematically using continuous dynamical variables. Their work concerns nerve excitation, not a complete theory of cognition, but it illustrates why a neuron cannot be understood only as an isolated on-or-off label. [2]

Binary representation does not restrict a computer to describing two-valued phenomena. A sequence of bits can represent a voltage, a probability, a differential equation, or an evolving model. The representation has finite precision, but finite precision alone does not establish that a required cognitive function is inaccessible. The relevant issue is how much precision, time, memory, and energy the function requires.

The same distinction matters for quantum computation. An n-qubit pure state can be represented by 2^n complex amplitudes. A classical numerical simulator can store and update those amplitudes for a finite quantum circuit with computable gate descriptions, approximating its outcome probabilities to a chosen accuracy. The direct method can be extremely expensive. Its existence nevertheless blocks a simple inference from quantum behavior to classical impossibility. Feynman's treatment of physical simulation helped establish why the resource cost of representing quantum systems matters. [3]

The 2^n state description is not a proof that every quantum computation requires exponential classical resources. Some circuits and some observable quantities admit much cheaper methods. A necessity argument has to identify the relevant computation and examine the best alternatives, rather than treat the cost of one simulation method as a universal lower bound.

There is also a gap between copying a brain and building intelligence. An artificial system could achieve the same functional result through a different mechanism. Establishing that the brain uses a quantum process would therefore answer a biological question first. Establishing that every successful AGI must use one would require a further argument.

My conjecture stands or falls on that further argument. The important question is whether a quantum resource supplies an indispensable capability at a cost that classical alternatives cannot match within the relevant limits.

4. What the research contributes

Decoherence and the timescale problem

Tegmark calculated decoherence times for particular neural and microtubule scenarios and argued that the relevant degrees of freedom should be treated classically. His estimates placed the loss of coherence far below the neural timescales he considered. This is a substantial challenge for proposals that require those quantum states to survive long enough to influence cognition. It is a model-based argument about specified mechanisms, not a measurement ruling out every possible quantum contribution to brain function. [4]

Hagan, Hameroff, and Tuszyński disputed whether those assumptions adequately represented the proposed microtubule mechanism. Their recalculation produced longer coherence estimates under different assumptions. The disagreement makes the experimental obligation clear: identify the physical state, measure its lifetime under relevant conditions, and demonstrate a functional consequence. Competing estimates do not establish that a brain performs the proposed quantum computation. [5]

Specific mechanisms are more useful than general appeals to quantum physics

Fisher proposed a possible route involving phosphorus nuclear spins and molecular structures that might preserve and distribute quantum information. The value of such a proposal is its specificity: it identifies candidate carriers, interactions, and links to neuronal activity. Those proposed links remain hypotheses in that paper. They provide possible experimental targets rather than evidence that general intelligence requires the mechanism. [6]

Babcock and colleagues combined theoretical analysis of collective ultraviolet excitation in tryptophan networks with measurements of fluorescence quantum yield in tubulin structures. Their results support an optical effect in organized biological material. The experiment does not demonstrate reasoning, learning, conscious experience, or an indispensable computational role in an intact brain. It is relevant to the physical possibilities being investigated, while leaving the cognitive claim open. [7]

A physical effect needs a functional explanation

Debates elsewhere in quantum biology provide a useful warning. Cao and colleagues reassessed claims about coherent energy transfer in photosynthesis and argued that long-lived spectral oscillations had been misinterpreted, with the relevant interexciton coherences too short-lived to perform the proposed function. A measured signal and its biological interpretation must be assessed separately. [8]

For AGI, the evidential chain is longer still: a physical effect must influence information processing, that influence must support a relevant capability, and that capability must resist adequate classical replacement under the stated resource limits. These papers address parts of the problem. They do not complete that chain. This is a focused discussion of relevant work, not a systematic review of every quantum-cognition proposal.

5. The quantum necessity conjecture

> **Conjecture: Human-level general intelligence, achieved within realistic physical resource limits, requires at least one functionally indispensable quantum information-processing resource. In an engineered system, that resource will operate alongside classical learning, memory, and control.**

To turn this into a particular scientific claim, define an evaluation standard T and a resource envelope R. T specifies the capabilities, test conditions, allowed assistance, and success criteria. R specifies limits on training and adaptation, inference time, memory, energy, and access to external computation. The operational claim is that a hybrid system can satisfy T within R and no admissible classical alternative can do so.

The word “no” is the hardest part of that statement. Experiments can compare selected systems. They cannot exhaust all possible algorithms or architectures. Persistent classical failure would support further investigation; it would not establish a universal impossibility result. A general necessity claim would need a lower bound, a physical constraint, or another argument extending beyond the tested implementations.

A brain-based argument also needs four distinct premises. First, a defined quantum mechanism operates under physiologically relevant conditions. Second, the mechanism makes a causal contribution to a specified cognitive capability. Third, that capability is necessary for the stated level of general intelligence. Fourth, no classical implementation supplies an adequate substitute within R. The conclusion depends on all four. Evidence for the first premise cannot stand in for evidence for the fourth.

Even if these premises held, quantum hardware would not automatically be sufficient for AGI. The system would still need a suitable architecture, learning process, memory, and interaction with its environment. Nor does standard quantum computation automatically supply a solution to undecidable problems. A claim requiring noncomputable physical dynamics would need a separate physical theory and separate evidence.

6. A research program that can reject the claim

Define success before examining the systems

An initial evaluation should cover six areas: learning an unfamiliar task, transfer between domains, causal reasoning under intervention, planning over extended sequences, retaining knowledge while

learning new material, and recognizing when available information is insufficient. These are proposed evaluation areas, not a claim that one fixed benchmark captures all of intelligence.

For each area, evaluators should use concealed tasks and declare a human comparison group with relevant preparation. A proposed threshold is performance at or above that group's median in every area, with uncertainty intervals reported and no severe deficit hidden by averaging scores. New task variants should test transfer rather than recall. The operational definition, task sampling procedure, failure thresholds, and resource budgets must be registered before the decisive comparison.

Make the biological intervention selective

A biological test should begin with a named mechanism and a measurable quantum signature. Researchers would then need an intervention that changes that signature while preserving, as far as possible, the ordinary biochemical and electrical functions that could otherwise explain a cognitive effect. Temperature, toxicity, structural damage, and changes in excitability are examples of competing explanations that require controls.

A reproducible loss of a specific capability following such an intervention, followed by a corresponding recovery when the mechanism is restored, would strengthen a causal account. It would not yet prove that an artificial classical system cannot reproduce the capability. That question belongs in the computational comparison.

Compare complete systems and credible replacements

A hybrid system should be compared with strong classical systems using the same task information and equivalent opportunities for tuning. Resource accounting should include data preparation, training, repeated measurement, error correction where used, communication, cooling, and the classical computation supporting the quantum component. A fast quantum subroutine can still belong to an inefficient overall system.

The critical experiment is replacement. Remove the proposed quantum resource and substitute the best available classical method, allowing the surrounding system to be retrained or redesigned within the declared budget. Merely switching off a module shows that the tested design depends on that module; it does not show that its quantum nature is indispensable. Classical alternatives should include materially different algorithms and, where relevant, physically realistic analog or neuromorphic approaches.

Independent groups should reproduce any claimed advantage on new tasks. The report should distinguish measured performance, extrapolation beyond tested sizes, and any theoretical assumption used to rule out classical alternatives. This proposal is not a completed preregistration; numerical budgets and task materials must be fixed before an actual experiment begins.

7. Predictions and disconfirming results

A useful prediction has to survive being stated in advance. The following commitments distinguish the conjecture from a general expectation that quantum technology will eventually be useful.

1. A consequential resource gap. If practical quantum necessity is correct, a relevant capability should exhibit a gap that remains consequential after complete resource accounting and serious classical optimization. A small speed improvement on an unrelated task would support only a narrower advantage claim.

2. A dependence that survives replacement attempts. For a system presented as evidence of necessity, an adequately resourced classical replacement should fail to preserve the required capability within the registered limits. If a classical replacement succeeds after redesign, that system no longer supplies the claimed evidence for necessity.

3. A measurable biological link, if the argument is biological. A theory locating the essential resource in the human brain should identify a reproducible quantum signature with a causal relationship to the claimed cognitive function. A failed result can reject that mechanism without rejecting every possible quantum route to engineered intelligence.

4. Acceptance of a classical counterexample. A demonstrably classical system that meets the preregistered standard T within R would refute the operational necessity claim for that standard and resource envelope. I would accept that result. The criteria cannot be changed afterward simply to preserve the prediction.

The interpretation of success also has to stay narrow. A hybrid system reaching the target first would establish a successful route. It would not prove that another route is impossible. Conversely, years of unsuccessful classical development would not establish quantum necessity: architecture, data, training, and evaluation could still explain the failures.

These are conditional scientific commitments, not a forecast of a particular discovery date. No percentage probability is assigned here. A more precise forecast would require a fixed target, deadline, and resolution procedure.

8. The objections that remain

The strongest objection is that the conjecture may add an unnecessary physical requirement. General intelligence could emerge from classical learning, memory, interaction, and sufficiently effective architectures. A classical system would not have to resemble a brain to demonstrate that possibility. The research program must therefore take classical success seriously, rather than explain it away as imitation.

A second objection is that cognitive tasks may not require access to the quantum details of their physical substrate. A coarse description can be sufficient for a particular function even when a finer description is physically more complete. Finding quantum behavior somewhere in a brain would not establish that a cognitive model must preserve it.

A third objection is that practical necessity changes with the resource envelope. A method that is unaffordable under one budget may become feasible under another. Any claim about what is practically necessary must preserve the original constraints when it is evaluated. Otherwise, a moving budget can turn a scientific statement into an unfalsifiable one.

These objections limit the strength of the conclusion. They also identify the work needed to strengthen it. The conjecture cannot be defended solely by pointing to the complexity of a brain, the failures of a particular model, or the unfamiliarity of quantum mechanics.

9. Conclusion

I expect quantum processing and classical AI to become deeply connected in attempts to achieve general intelligence. My stronger conjecture is that a functional quantum resource will prove necessary for human-level general intelligence within realistic physical limits. That is the position I am putting forward.

The case has to be made at the level of mechanism and capability. The brain's biological complexity does not establish a limit on binary computation, and a quantum effect does not establish a requirement for AGI. What would matter is a demonstrated connection between a quantum resource, an indispensable cognitive function, and the failure of adequate classical substitutes under explicit constraints.

If a classical system meets the agreed standard, the operational claim fails. If a quantum resource proves indispensable under that standard, the claim gains support. Either result would teach us something substantial about intelligence. The purpose of stating the conjecture now is to make that future judgment possible.

References

[1] Morris, M. R., Sohl-Dickstein, J., Fiedel, N., Warkentin, T., Dafoe, A., Faust, A., Farabet, C., and Legg, S. (2023; revised 2025). Levels of AGI for Operationalizing Progress on the Path to AGI. arXiv:2311.02462, version 5. https://arxiv.org/abs/2311.02462v5

[2] Hodgkin, A. L., and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. The Journal of Physiology, 117(4), 500–544. https://doi.org/10.1113/jphysiol.1952.sp004764

[3] Feynman, R. P. (1982). Simulating physics with computers. International Journal of Theoretical Physics, 21, 467–488. https://doi.org/10.1007/BF02650179

[4] Tegmark, M. (2000). Importance of quantum decoherence in brain processes. Physical Review E, 61, 4194. https://doi.org/10.1103/PhysRevE.61.4194

[5] Hagan, S., Hameroff, S. R., and Tuszyński, J. A. (2002). Quantum computation in brain microtubules: Decoherence and biological feasibility. Physical Review E, 65, 061901. https://doi.org/10.1103/PhysRevE.65.061901

[6] Fisher, M. P. A. (2015). Quantum cognition: The possibility of processing with nuclear spins in the brain. Annals of Physics, 362, 593–602. https://doi.org/10.1016/j.aop.2015.08.020

[7] Babcock, N. S., Montes-Cabrera, G., Oberhofer, K. E., Chergui, M., Celardo, G. L., and Kurian, P. (2024). Ultraviolet Superradiance from Mega-Networks of Tryptophan in Biological Architectures. The Journal of Physical Chemistry B, 128, 4035–4046. https://doi.org/10.1021/acs.jpcb.3c07936

[8] Cao, J., et al. (2020). Quantum biology revisited. Science Advances, 6(14), eaaz4888. https://doi.org/10.1126/sciadv.aaz4888