

Article

Not peer-reviewed version

Spacetime Bounds on Consciousness

Michael Timothy Bennett *

Posted Date: 13 March 2026

doi: 10.20944/preprints202602.1708.v2

Keywords: consciousness; cognition; stack theory; unity; conscious binding



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Spacetime Bounds on Consciousness

Michael Timothy Bennett

School of Computing, The Australian National University, Canberra, ACT 2601, Australia; michael.bennett@anu.edu.au

Abstract

Is an ant colony conscious? What about a group of people talking, a cloud-hosted language model, or even a galaxy? Can a conscious mind only get so big? Does consciousness depend only on what is computed, or when and where? I see two possibilities affecting the answers to these questions. I name them Chord and Arpeggio, and formalize the distinction mathematically. If the ingredients of a subjective experience must be simultaneously true at one objective instant and causally exchange influence within a time window θ , then the system diameter D satisfies $D \leq \kappa v \theta$, where v is the signal speed ceiling and κ depends on exchange architecture. I call this requirement Chord, because it is like a musical chord whose notes sound together. The alternative is Arpeggio. It asks only that each ingredient occur somewhere in the window. I prove that Arpeggio is strictly weaker than Chord, and that architectures with limited concurrency can satisfy Arpeggio while structurally forbidding Chord. I argue for Chord on formal, neuroscientific, and architectural grounds. A mechanistic model confirms a fragmentation transition at the theoretical threshold. I examine primate corpus callosum data to estimate empirical lower bounds on θ . I provide case studies showing that under Chord, ant colonies and human populations are ruled out as single conscious entities, cloud-hosted AI is constrained by co-instantiation rather than diameter, and brain-computer interface hybrids face latency-dependent limits. A mind can only get so big. Arpeggio is far more permissive, implying consciousness seemingly everywhere.

Keywords: consciousness; cognition; stack theory; unity; conscious binding

Introduction

A conscious moment feels unified. An instant of time. Yet brains are spatially extended and coordination across space *takes* time. This raises the question of how an experienced “instant” works in relation to the objective progression of physical events.

To deal with this, many accounts of consciousness posit a temporal integration window of duration θ . Within this window the system is allowed to gather and bind ingredient facts such as sense, feeling, and thought, to which one has reportable conscious access [1]. An ingredient is a grounded physical fact, meaning a property of the underlying microstate rather than a purely internal label. Microstate here means the full physical state at the resolution where the ingredients live. At the neuronal scale, for example, it could mean the full pattern of spikes, membrane potentials, and synaptic states. This windowing idea is widespread. For example it appears in global workspace accounts [2,3], recurrent processing theories [4], dynamical synchrony [5–7], integrated information [8] and stack theoretic proposals [9–12]. In Stack Theory, meaning is ontological. The meaning of a statement is its truth set, which is the set of microstates that satisfy it. It also appears in classic philosophical discussions of subjective timing [13]. The interdependence of temporal and spatial dimensions of perceptual experience has also been emphasised from a phenomenological perspective [14]. Empirically, psychophysics supports integration on the order of tens to hundreds of milliseconds, with task-dependent structure [15–18]. The obvious question is what fixes θ . This paper does not treat θ as a magical constant. For a candidate content ℓ , any admissible window must be long enough to permit the required reciprocal exchange and, if persistence is the mechanism used to close the

Temporal Gap, short enough that the grounded ingredients still belong to one moment. The lower and upper constraints are derived explicitly below.

A window may be necessary, but it is not sufficient for literal co-presence [12]. Each ingredient can occur somewhere inside the window even when there is no instant at which they are all true together in the underlying physical state. I call that mismatch the Temporal Gap. In this paper, Arpeggio means every ingredient becomes true at least once during the window. Chord means there exists an instant during the same window when all ingredients are true together. Chord and Arpeggio are musical terms. A chord is a set of notes played at the same time, while an arpeggio is the notes of a chord but played in sequence. The notes of a musical chord resonate and act upon one another, creating harmonics and phasing. The notes of an arpeggio do not. In what follows, I treat Chord as the requirement for a unified, integrated subjective moment of experience in the strong sense studied here. Arpeggio is the weaker alternative. It only requires that each ingredient appears at least once somewhere in the window. I then provide arguments in support of Chord (extending a workshop paper [12] on which this paper is to some extent based) and case studies for Chord versus Arpeggio in liquid versus solid brains [19], brain-computer interface hybrids, cloud-hosted AI, human populations, and the limiting case of Arpeggio applied without constraint, as well as human-AI hybrids [11,20]. Formal definitions appear in Definitions 2 and 3.

1. I extend a workshop paper [12] that introduced the Stack Theory time semantics of Arpeggio and Chord. Here I add a within-window causal exchange postulate to Chord and derive a spacetime diameter bound for one unified moment.
2. I integrate the algebraic results from that workshop paper (the non-commutation of existential temporal lift with conjunction, the concurrency capacity threshold, and the persistence collapse condition) and give arguments in favour of Chord over Arpeggio on formal, neuroscientific, and architectural grounds.
3. I build a falsification pipeline, validate it with stress tests, and anchor the lower side of the window using a reanalysis of published primate corpus callosum microstructure.
4. I address the “make the window larger” objection. The window is bracketed from below by the exchange budget implied by D , κ , and v , and from above by ingredient persistence when persistence is the mechanism invoked to close the Temporal Gap. This yields a feasibility interval rather than an arbitrary stipulation.
5. I present case studies for liquid brains, brain-computer interface hybrids, cloud-hosted AI, human populations, and Arpeggio taken to its logical extreme.

Where does θ come from? In this paper θ is not without limits. It is the horizon over which the system’s binding mechanism can keep the relevant ingredients available for mutual influence. In biological brains, candidate mechanisms include recurrent loops, synaptic integration, oscillatory coherence, and ignition-like dynamics [3,4,7]. Each carries an intrinsic timescale, so θ is expected to lie in a constrained range rather than be arbitrary. Empirically, θ can be estimated from psychophysics, from neural markers of integration such as phase-locking or ignition, and from task-dependent report latencies. Crucially, enlarging θ without changing the ingredient set is limited by ingredient stability. If an ingredient typically turns over on timescale τ_p , a much longer window either reopens the Temporal Gap or forces a coarser redescription of what the ingredients are. The section on window constraints makes this tradeoff explicit.

Everything that follows is conditional. The bound only applies in the case of the Chord postulate. If co-instantiation of ingredients and within-window exchange are not required, then the diameter bound does not apply.

Notation preview. A candidate moment is a time window $[t, t + \theta]$. D is the diameter of the grounded support, meaning the largest distance across the physical sites that carry the grounded ingredients. v is an upper bound on causal propagation speed in the relevant substrate. κ is an architecture factor determined by which sites must exchange messages within one window.

Box 1: Chord requirements and bound.

Assume a finite grounded support for a conscious moment, meaning a finite set of physical sites whose grounded ingredients participate in that moment:

1. **Co-instantiation.** A unified moment contains at least one instant inside its window where all grounded ingredients are simultaneously true (meaning they are all part of one physical microstate).
2. **Within-window exchange.** During the same window, grounded ingredients exchange influence by two-way communication along an exchange graph, and signals propagate no faster than v .

Then. If grounded support (physical system) has diameter D and window is θ , then $D \leq \kappa v \theta$. To get a big unified moment, you buy it with time, speed, or architecture.

Supplementary Note 0 in the SI is a glossary and figure reading guide for the technical terms and plots. If you are coming from another discipline, start there.

Temporal integration windows are not new. The novelty here is turning the Chord requirements of co-instantiation plus within-window exchange into a falsifiable spatial inequality with an explicit architecture factor. In other words, if a moment must be causally knitted together within its own window, then unity has a speed limit. That speed limit converts into a size limit. To put it provocatively, a conscious mind can only be so big for a given signal speed.

Results

Windowing Is Weaker Than Simultaneity

I represent the content of a candidate moment using a Stack Theoretic abstraction layer [21]. An abstraction layer is built from an environment Φ and a finite vocabulary \mathfrak{v} . I define these objects next. The key modelling choice is extensional semantics. It treats each ingredient as a physical constraint and it treats meaning as the set of microstates that satisfy that constraint.

Definition 1 (Environment, programs, statements, and truth sets). *An environment is a nonempty set Φ of mutually exclusive microstates. For any set X , write 2^X for its power set, meaning the set of all subsets of X . A program is any set $p \subseteq \Phi$. Write $\mathcal{P} := 2^\Phi$ for the set of all programs. A vocabulary $\mathfrak{v} \subseteq \mathcal{P}$ is a finite set of programs that a system can implement. The induced embodied language is*

$$L_{\mathfrak{v}} := \{ \ell \subseteq \mathfrak{v} \mid \bigcap_{p \in \ell} p \neq \emptyset \}.$$

Elements $\ell \in L_{\mathfrak{v}}$ are called statements. The truth set of a statement is

$$T(\ell) := \bigcap_{p \in \ell} p \subseteq \Phi, \quad T(\emptyset) := \Phi.$$

Narrative interpretation. A microstate is one fully specified physical configuration. A program is a yes or no property of microstates, represented by the set of microstates where it holds. In Stack Theory, the word program does not mean executable code. It means a physically implementable predicate. The notation 2^Φ means all subsets of Φ . It is the set of all possible programs you could define. A vocabulary is the particular finite menu of programs this system can actually implement. A statement ℓ is a bundle of programs that can all be true together. Because the vocabulary is finite, every statement is automatically a finite bundle. Its truth set $T(\ell)$ is the set of physical states where the whole bundle is true at once. In this paper an ingredient is one program $p \in \ell$. The vocabulary is finite because an embodied system has finite information capacity [22,23]. The Bekenstein bound is one concrete

physics limit of this kind. It upper bounds how much information a finite region can store. You do not need to accept the Bekenstein bound in particular. Any finite physical memory implies a finite menu of physically distinguishable conditions at a chosen resolution. That is why treating vocabularies as finite is a reasonable modelling default. In Stack Theory, the pair (Φ, ν) plus its induced language is the abstraction layer used in this manuscript. Supplementary Note 2 restates the Stack Theory definitions used here.

Temporal binding introduces a window of duration θ . Let $\phi(\tau) \in \Phi$ denote the underlying microstate at objective time τ . I call a program grounded when it is evaluated directly on $\phi(\tau)$ in the base environment Φ , rather than inside a separate internal label space [12]. Supplementary Note 7 restates the exact discrete-time version used in Appendix 26. The main text keeps the same logical distinction but writes it in continuous-time interval notation.

Definition 2 (Arpeggio and co-instantiation). *Fix a statement $\ell \in L_\nu$ and a time window $\sigma = [t, t + \theta]$. Ingredient-wise occurrence holds when each ingredient becomes true at least once somewhere inside the window.*

$$\text{Occur}(\ell, \sigma) \iff \forall p \in \ell \exists \tau \in \sigma \text{ such that } \phi(\tau) \in p.$$

Co-instantiation holds when there exists at least one instant inside the window at which all grounded ingredients are simultaneously true.

$$\text{CoInst}(\ell, \sigma) \iff \exists \tau \in \sigma \text{ such that } \phi(\tau) \in T(\ell).$$

Definition 3 (Temporal Gap). *Fix a statement $\ell \in L_\nu$ and a time window $\sigma = [t, t + \theta]$. The Temporal Gap event is when every ingredient occurs somewhere inside the window, but there is no instant inside the window when they are all true together.*

$$\text{TempGap}(\ell, \sigma) \iff \text{Occur}(\ell, \sigma) \text{ and } \neg \text{CoInst}(\ell, \sigma).$$

Narrative interpretation. The symbol σ is just shorthand for the window $[t, t + \theta]$. Arpeggio says each ingredient shows up at least once somewhere inside the window. Chord says there is at least one instant when the ingredients overlap in the underlying physical state. Chord automatically implies Arpeggio. If all ingredients are true together at one instant, then each ingredient occurs at that instant. Arpeggio is therefore a weaker criterion than Chord. The Temporal Gap is the mismatch where Arpeggio holds but Chord fails. It is illustrated in Figure 1.




Figure 1. Both occur in the window, but never at the same instant. This is the Temporal Gap. Arpeggio can hold without co-instantiation.

This gap matters because if consciousness requires an instant at which the ingredients are all grounded together, then the theory needs a synchrony condition, not just a set of events that happen separately within a shared clock interval.

Diameter Bound

I add a within-window causal exchange postulate. If ingredients are jointly instantiated, then the sites that host those ingredients must be able to exchange causal influence within the same window.

Assumption 1 (Within-window causal exchange). Let (X, ρ) be the physical substrate equipped with a metric ρ . Let $S \subseteq X$ be a finite grounded support for a moment, meaning a finite set of physical sites whose grounded ingredients participate in that moment. Define its diameter by

$$D := \max_{u, u' \in S} \rho(u, u').$$

Assume there exists a connected undirected exchange graph $G = (S, E)$. For every required edge $\{u, u'\} \in E$ there are causal travel times $\eta_{u \rightarrow u'}, \eta_{u' \rightarrow u} \geq 0$ such that a round trip fits inside the window

$$\eta_{u \rightarrow u'} + \eta_{u' \rightarrow u} \leq \theta.$$

Signals propagate no faster than v . Equivalently, each one-way travel time must satisfy $\rho(u, u') \leq v \eta_{u \rightarrow u'}$ and $\rho(u, u') \leq v \eta_{u' \rightarrow u}$.

Narrative interpretation. v is a speed limit (not to be confused with the embodied vocabulary symbol v). A metric is just a rule for distance that behaves like ordinary distance and obeys the triangle inequality. Triangle inequality means going from a to c directly is never longer than going from a to b to c . Connected means you can reach any site from any other site by hopping along required exchange edges. If the exchange graph is not connected, your candidate moment is already fragmented into components, so apply the bound component by component. Pick the physical sites that are supposed to contribute grounded ingredients to the moment. A site can be a neuron, a cortical column, a brain region, or any other spatial unit that makes sense for the modelling scale. Draw an edge between two sites if your theory says those two sites must be able to influence each other within one window. The travel times $\eta_{u \rightarrow u'}$ and $\eta_{u' \rightarrow u}$ are the message times in each direction. The inequality $\eta_{u \rightarrow u'} + \eta_{u' \rightarrow u} \leq \theta$ says there is time for a signal to go out and for an acknowledgement to come back within the same moment. The speed bound v says these travel times cannot be arbitrarily small because no signal can outrun v . Together with the round trip budget, this implies that every required edge can span at most $v\theta/2$ of physical distance.

Pick the distance rule ρ and the speed bound v as a matched pair. If you measure ρ as straight line distance, then v should be an effective speed that already includes wiring detours. If you measure ρ as path length along fibres or wires, then v can be a more direct conduction speed along that path.

The exchange graph specifies who has to talk to whom within one subjective moment. The round trip condition is the minimal way of cashing out the idea that a moment is causally stitched together rather than merely co-timed.

In the limits the architecture can be a hub, an all-to-all mesh, or a general graph. Each architecture determines a geometric constant κ that converts time into diameter. The formal graph definition of κ and the proof are in Supplementary Note 3.

Theorem 1 (Spacetime diameter bound). Under Assumption 1, the grounded support diameter satisfies

$$D \leq \kappa v \theta.$$

For hub exchange, meaning a star graph, $\kappa = 1$. For all-to-all exchange, meaning a complete graph, $\kappa = \frac{1}{2}$. In general, let $d_G(u, u')$ be the length in edges of a shortest path between u and u' in G . Define the hop diameter

$$h(G) := \max_{u, u' \in S} d_G(u, u').$$

Then $\kappa = h(G)/2$.

Narrative interpretation. A star graph is a hub and spokes. A complete graph connects every pair directly. The hop distance $d_G(u, u')$ counts how many required handoffs a message must make to get from u to u' . The hop diameter $h(G)$ is the worst case hop count across the whole support. The

constant $\kappa = h(G)/2$ is therefore a pure architecture factor. Each hop can span at most $v\theta/2$ of physical distance because a round trip must fit inside the window. So the bound says the diameter is at most $h(G)$ hops worth of distance, which is exactly $\kappa v\theta$.

Units reminder. D is a distance. v is distance per time. θ is time. So $v\theta$ is a distance and κ is dimensionless.

Worked example. Suppose $\theta = 20$ ms and $v = 10$ m s⁻¹. Then $v\theta = 0.2$ m. Under hub exchange, meaning $\kappa = 1$, the bound gives $D \leq 0.2$ m. Under all-to-all exchange, meaning $\kappa = \frac{1}{2}$, the bound gives $D \leq 0.1$ m. These are order-of-magnitude numbers, the point of them being that a time budget becomes a size budget.

In other words, if the window is too short for causal exchange across the support, there is no way to knit the ingredients into one co-instantiated state. See Figure 2 for a basic illustration.




Figure 2. Geometry underlying the bound. It is a worst case. In a star architecture, the farthest pair of contributors communicates in two hops via the hub, so each hop must fit inside half the window and the overall diameter budget is $D \leq v\theta$. In a complete architecture, the farthest pair must complete a round trip over distance D within the window, giving the tighter budget $D \leq v\theta/2$.

Fragmentation Transition

Theorem 1 is a worst case bound. To see whether the same geometry shows up in a concrete mechanism, I built a minimal message-passing model.

Assume there are N sites arranged in a circle of diameter D . A moment is a window of duration θ which is how much time we have to integrate the ingredients before they degrade in some manner. Signals propagate at speed at most v , so a one-way message over distance d has baseline travel time d/v . On top of this baseline each directed message receives an independent exponential jitter term with mean $j\theta$. Exponential here means most messages are close to their baseline travel time but a few run very late. It is a simple way to model rare but consequential delays. Here j is dimensionless and can be read as the expected fraction of the window lost to random delays.

I evaluate two exchange graphs. Hub exchange means a star graph. All-to-all exchange means a complete graph. The model declares the moment integrated only if every required edge completes a two-way exchange within the same window. This yields an exact success probability. Success probability is the chance, over the random jitter, that the window meets this integration requirement. The full derivation is given in Supplementary Note 4 [24]. To present results compactly, define the dimensionless size ratio

$$x := \frac{D}{v\theta}.$$

This ratio is the fraction of the window needed to traverse the diameter at speed v .

Narrative interpretation. x is a score of how close you're cutting it time-wise (analogous to asking if you are going to be 5 minutes early to the job interview, or 5 seconds). It compares the best case time to cross the support once, D/v , to the time you have, θ . Small x means plenty of slack, while large x means you are spending most of the window just moving signals around (analogous to running to your job interview because you cut it too fine). A run of the mechanistic model counts as a success only if every required two-way exchange finishes before the window ends. Figure 3a plots this success probability as a function of x . The curves stay close to 1 for a while and then drop rapidly toward 0.

That rapid drop is what I mean by collapse here. It is the point where the window becomes too tight for within-window exchange. Hub exchange drops near $x \approx 1$. All-to-all exchange drops near $x \approx \frac{1}{2}$.

Figure 3b is a robustness sweep. A sweep means rerunning the same model many times while changing input parameters. Here I vary two nuisance parameters, the number of sites N and the jitter level j . j is the mean random extra delay per message written as a fraction of the window. For each (N, j) setting I summarise the transition by x_{50} , the value of x where the success probability is 0.5. Across $N \in \{10, 20, 40, 80\}$ and $j \in \{0, 0.005, 0.01, 0.02\}$, with $j = 0$ interpreted as the deterministic no-jitter limit, I obtain

$$x_{50}^{\text{hub}} \in [0.864, 1.000], \quad x_{50}^{\text{all}} \in [0.413, 0.500].$$

The key point is that x_{50} stays close to the theoretical thresholds. This suggests the transition is not a fragile artifact of one particular choice of N and j .




Figure 3. Mechanistic integration model. Panel (a) shows the probability that a window is successful, meaning every required two-way exchange finishes before the window ends. The horizontal axis is $x = D/(v\theta)$, the ratio of best-case diameter-crossing time to the available window. Panel (b) is a robustness sweep. It repeats the same calculation while varying the number of sites N and the jitter level j . Here j is the mean random extra delay per message as a fraction of the window. It reports x_{50} , the value of x where success probability is 0.5. Reference lines mark the theoretical thresholds $x = 1$ for hub exchange and $x = \frac{1}{2}$ for all-to-all exchange.

Primate Scaling Anchor

I wanted to know how this bound might manifest in the real world. To get a rough empirical anchor, I use an existing primate dataset. It is already published and requires no new experiments.

The corpus callosum is the major fibre tract connecting the two cerebral hemispheres. Any moment that requires causal exchange between hemispheres cannot complete faster than the relevant callosal conduction time.

Phillips et al measured callosal axon diameter distributions across primates and reported implied one-way interhemispheric conduction times for two fibre proxies [25,26]. The corrected Table 1 contains $n = 15$ individual primates across 14 primate species. The scaling fits below are descriptive. They are not a phylogeny-aware comparative analysis that corrects for shared ancestry. One proxy uses the median axon diameter. The other uses the 95th percentile axon diameter as a fast fibre proxy. I use these implied conduction times as reported in the corrected Table 1, rather than recomputing them. These one-way delays are floors on any integration window that needs information to cross between hemispheres. Under the stronger two-way exchange postulate used elsewhere in this paper, a round trip constraint would be roughly twice these values. The scaling pattern is unchanged because the factor is constant. This does not claim that callosal delay alone determines consciousness. It only supplies a physically grounded timescale that any bilateral moment must at least accommodate. It is an illustration using real-world data. Interhemispheric conduction delay has long been proposed as a

constraint on integration and hemispheric specialisation, and it scales with brain size [27–29]. Figure 4a plots the implied lower bound θ_{\min} versus brain mass. On a log-log fit, the median axon proxy scales with exponent 0.224 with bootstrap interval [0.168, 0.283] and $R^2 = 0.84$. A bootstrap interval is an uncertainty interval computed by resampling the data. R^2 is the usual fraction of variance explained by the fit. Supplementary Note 6 describes the regression and bootstrap procedure used to compute these values. It also explains how to read the primate plots. The fast axon proxy scales with exponent 0.197 with interval [0.116, 0.289] and $R^2 = 0.71$. Across the sample, brain mass spans a factor 148 while θ_{\min} spans factors 3.36 and 3.12 for median and fast proxies.

Figure 4b converts the same data into a direct constraint. For any candidate window θ , it reports the fraction of individuals whose reported conduction time exceeds θ . For example, 0.267 of median proxy times exceed 25 ms, while 0.667 of fast proxy times exceed 5 ms.

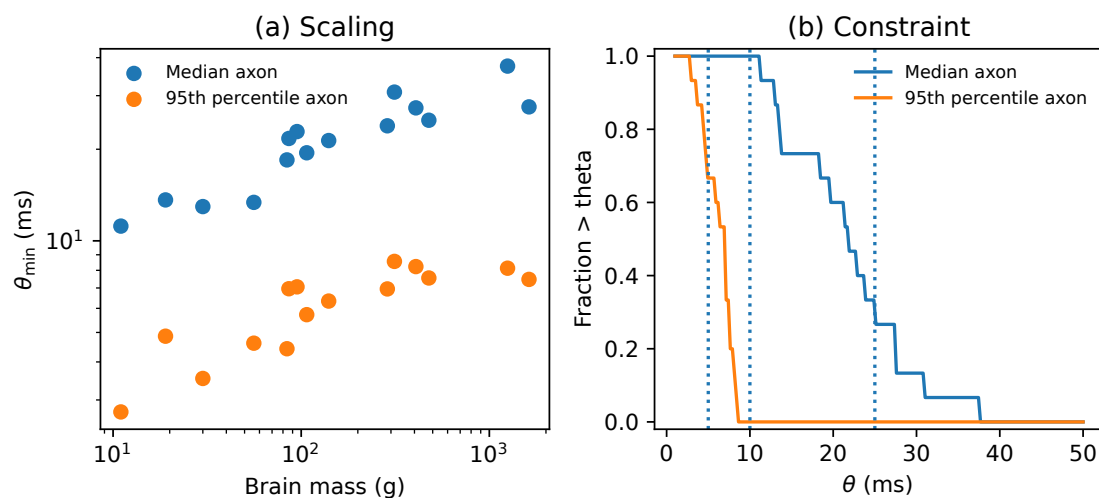


Figure 4. Primate literature anchor. Dataset contains $n = 15$ individuals from the corrected Phillips et al Table 1 [26]. Panel a shows the implied lower bound θ_{\min} from published interhemispheric conduction time proxies (median and fast fibres) versus brain mass. Panel b shows, for a candidate window θ , the fraction of individuals whose reported conduction time would exceed that window.

Falsification and Stress Tests

Claims like $D \leq \kappa v \theta$ are only scientifically useful if they can be falsified. Given data, the natural falsification target is the margin

$$M := \hat{D} - \kappa \hat{v} \hat{\theta}.$$

Here \hat{D} is a diameter estimate for the candidate support, $\hat{\theta}$ is the candidate window duration, and \hat{v} is an estimate of the fastest relevant causal propagation speed. A violation corresponds to $M > 0$.

Testing protocol. To test this bound:

1. Pick a candidate unity marker and a window duration θ .
2. For each window, estimate the support diameter \hat{D} and collect fast-propagation samples to estimate \hat{v} .
3. Compute the margin $M = \hat{D} - \kappa \hat{v} \hat{\theta}$.
4. If windows overlap heavily, thin them by keeping every k th window.
5. Test whether the mean margin is greater than zero using a one sided t test.

The fiddly part is \hat{v} . If you underestimate the true speed ceiling, you can manufacture violations. To stay conservative, I treat v as a sampled quantity and estimate \hat{v} using a high quantile, specifically the 95th percentile of the within-window samples. The 95th percentile is the value below which 95% of the samples fall. Call this the q95 rule. q95 ignores the slow bulk and asks how fast the system can plausibly go.

I then test whether the mean margin across windows is positive using a one sided t test. This asks whether the average violation is reliably above zero, given sampling noise. Windows can be statistically dependent. To avoid overstating certainty, I control dependence by thinning. Thinning means keeping only every k th window so that the remaining windows are closer to independent. For example, if you slide a 50 ms window forward in 1 ms steps, then adjacent windows share almost all of their data. Keeping every 50th window removes that overlap. Thinning pretty much means making the overlap smaller. Supplementary Note 5 gives the full protocol. It also explains how to read each stress test plot. I validate the full pipeline using Monte Carlo based stress tests. I simulate windowed data where the true ratio $\frac{D}{\kappa v \theta}$ is known, run the same falsification procedure, and measure false refutations and power. Power here means the probability of detecting a true violation. The stress tests support three takeaways.

1. The false refutation rate is controlled at or below the nominal level when v is conservatively estimated. Under the default q95 rule used here I observe no false refutations at ratio 1.00 in 2500 replications.
2. Power rises rapidly once the true ratio exceeds one. In these Monte Carlo tests I set $\kappa = 1$ as a units choice, so the x axis should be read as $\frac{D}{v \theta}$ in general. With the conservative default quantile estimator used here and $n = 40$ windows, the rise begins once the true ratio exceeds one by only a few percent.
3. Naive estimators that underestimate the relevant speed bound, especially low end choices like taking the minimum within-window sample, can manufacture violations and inflate false refutation dramatically.

All simulation code and source data are included in the publicly available supplementary information [24].

Discussion

The inequality $D \leq \kappa v \theta$ is a conditional claim about what a physically unified moment requires. It follows from two modelling commitments. The first commitment is the Chord requirement [12]. Unity requires not just that ingredients occur somewhere in a window, but that there exists an instant of objective time when the grounded ingredients are jointly true. This rules out the Temporal Gap. Stack Theory formalises this by making the joint truth condition membership in a joint truth set [12]. The second commitment is within-window causal exchange. If the ingredients are supposed to be one moment rather than a list, then they must be able to constrain each other by two-way communication within the same window. This turns a time budget into a spatial budget. If either commitment is false for the relevant notion of consciousness, then the bound does not apply as intended.

My mechanistic model shows that even a toy integration mechanism exhibits a sharp fragmentation transition near the theoretical thresholds. The primate reanalysis is an empirical sanity check on scale. By this I mean the callosal conduction time is a concrete example of latency, one that any bilateral moment must at least accommodate. The lower bound scales with brain size, so it remains consistent with the idea that larger brains either integrate more slowly, integrate more locally, or change architecture in a way that reduces κ .

Why Chord and not Arpeggio?

Chord requires two things of a unified conscious moment. First, the grounded ingredients must be co-instantiated at some instant of objective time within the window. Second, the sites hosting those ingredients must complete within-window causal exchange. Arpeggio requires only that each ingredient occurs somewhere in the window. Neither co-instantiation nor causal exchange is demanded. Why prefer Chord?

I give arguments for each requirement in turn. More detailed formal treatment appears in Supplementary Notes 7 and 8, which integrate the algebraic framework developed in prior work [12].
Evidence for co-instantiation.

The formal motivation starts with a quantifier order result. In prior work I proved that the existential temporal lift \diamond_{Δ} does not commute with conjunction [12]. Concretely, in the appendix-compatible discrete notation with window $\sigma = (\phi_0, \dots, \phi_{\Delta})$, ingredient-wise occurrence asks $\forall p \in \ell \exists k \leq \Delta$ such that $\phi_k \in p$. Co-instantiation asks $\exists k \leq \Delta$ such that $\forall p \in \ell, \phi_k \in p$. These differ by a quantifier swap. For a window of more than one instant and a statement with more than one ingredient, ingredient-wise occurrence is strictly weaker than co-instantiation (Supplementary Note 7, restating Theorem 3 of [12]).

In musical terms, this means a system can play every note at *different* times, without there ever being a moment when all the notes are played *together*. Should consciousness require the joint truth of its grounded ingredients at some objective instant, then Arpeggio is too weak to guarantee it. Concurrency capacity is a well-defined measure of a system's capacity to play a chord. The greater the concurrency capacity, the larger the chord it can play. An architecture that can activate at most c contributors at a time can satisfy ingredient-wise occurrence for an n -ingredient conjunction (with $n > c$) by cycling through contributors across the window. But it can never co-instantiate all n contributors because no single instant has more than c active (Supplementary Note 8, restating Theorem 5 of [12]). Under Arpeggio, a strictly sequential processor that cycles through features one at a time would count as hosting a unified moment. Under the Chord postulate, it would not.

Neuroscience provides tentative empirical support for the co-instantiation requirement. In masking paradigms, conscious perception correlates with transient long-range gamma phase synchrony across widely separated cortical areas, even when local gamma power is similar for seen and unseen stimuli [30]. What distinguishes conscious from unconscious processing is not the occurrence of local activity for each feature, but the transient episode of large-scale temporal coordination in which distributed areas synchronise together. This is closer to co-instantiation than to ingredient-wise occurrence.

During non-REM sleep, TMS-EEG responses become strong locally but fail to propagate, consistent with a breakdown of effective connectivity when consciousness fades [31]. Each area can still respond in isolation, but they do not respond together. Arpeggio-like ingredient-wise activity persists. What is lost is the coordinated, simultaneous engagement that Chord demands.

More broadly, perturbational complexity measures like PCI track the level of consciousness across waking, sleep, anaesthesia, and disorders of consciousness [32]. These measures quantify the extent to which a perturbation spreads through the brain in a differentiated yet integrated manner. They are sensitive not to whether regions respond at all, but to whether they respond as a coordinated whole. This pattern is consistent with the idea that the system-level regime associated with consciousness is one of temporally coordinated co-engagement, not merely sequential activation of parts.

There is also a persistence argument (Supplementary Note 7). If ingredients, once true, remain true for the rest of the window, then ingredient-wise occurrence automatically entails co-instantiation. Persistence closes the Temporal Gap. Without persistence, the gap can remain open. Whether biological substrates exhibit sufficient within-window persistence for the relevant grounded ingredients is an empirical question, but the formal point is that the gap between Arpeggio and Chord is closed only under a nontrivial dynamical condition that the substrate must actively satisfy.

None of these lines of evidence prove that co-instantiation is necessary for consciousness. They do, however, make the permissiveness of Arpeggio less compelling. If conscious level covaries with coordinated, simultaneous neural dynamics rather than with mere ingredient-wise activation, then Chord is the more parsimonious postulate.

Evidence for causal exchange. Suppose co-instantiation matters. A natural follow up question is why it would matter. Co-instantiation alone guarantees that the ingredients are jointly true at one instant, but it does not guarantee that they constrain each other. Two distant neurons could happen to be in their respective target states at the same time by coincidence, without either one's state being causally influenced by the other. If such a coincidence counted as a unified moment, then conscious unity would not require physical integration, only temporal alignment. If only temporal alignment were

required, how could we justify that claim? Why would two things need to exist at the same time if they don't interact?

Returning to the musical analogy, think of how the notes of a musical chord interact. They resonate and affect the environment in which they are played, and their interaction produces harmonics. The whole is something very different from the parts, both to the listener's ear and in the underlying physical consequences.

The causal exchange postulate addresses this. It says that the co-instantiated ingredients must also exchange influence within the same window. The sites hosting the ingredients must be able to send and receive signals from each other before the window closes. This is a minimal way to formalise the idea that a moment is causally stitched together rather than merely co-timed.

Recurrent processing theory provides a concrete neural precedent. Lamme argues that feedforward processing alone is insufficient for conscious perception, and that recurrent exchange between areas is needed [4]. The critical step is not just that signals arrive at multiple areas within a time window, but that those areas engage in two-way communication within that window. This is architecturally the same structure as the within-window causal exchange postulate.

Global workspace theory points in a similar direction. The workspace model posits that information becomes conscious when it is broadcast widely and made available to multiple specialised processors simultaneously [2,3]. Broadcasting is a causal process. It requires that distributed modules receive and are influenced by shared information within the relevant time window. Arpeggio would allow the ingredients to activate in sequence without any mutual constraint. Chord with causal exchange captures the workspace intuition that the processors must be simultaneously engaged and mutually informed.

The communication through coherence hypothesis is also relevant [7]. It holds that effective neural communication depends on phase alignment between oscillating neural populations. When populations are phase-aligned, spikes from one population arrive during the excitable phase of another, enabling mutual influence. When they are not aligned, the same signals fail to drive the target population. This is a biologically specific version of within-window causal exchange. The "window" is set by the oscillation period, and "exchange" is the mutual driving that occurs during phase alignment.

In summary, co-instantiation ensures the ingredients are jointly true at one objective instant. Causal exchange ensures they actually constrain each other within the same window. Together they constitute the Chord postulate. Dropping either one allows systems that most theories of consciousness would not recognise as unified. The bound $D \leq \kappa v \theta$ follows from the conjunction of both requirements.

Novelty in Comparison to a Similar Bound

There is conceptual overlap with a persistence condition proposed by Sendall [33]. However the bound I present here differs in three key respects. First, my architecture factor $\kappa = h(G)/2$ resolves the dependence on exchange topology that Sendall's effective causal diameter leaves implicit. Second, substituting the substrate-specific signal speed v for the speed of light yields bounds that are actually constraining for biological and engineered systems, not just for exotic spacetimes. Third, the mechanistic fragmentation model, the primate empirical anchoring, and the falsification protocol give the bound testable content. Sendall's horizon result is an interesting limiting case of the present framework, and so can be understood as complementary. At a strict one-way boundary the effective round-trip speed is zero, so $D \leq \kappa \cdot 0 \cdot \theta = 0$ and no support of positive diameter can straddle the boundary.

What Constrains the Window?

One obvious objection is that the bound $D \leq \kappa v \theta$ can always be satisfied by making θ larger. That would be fatal if θ were free. It is not.

First, θ is a mechanistic parameter, not a convention. In most neuroscientific models that use integration windows, the window is set by a process such as recurrent loops, ignition-like thresholding, phase-coherence cycles, synaptic integration, or related dynamical motifs [3,4,7]. Those processes

have intrinsic time constants. For example, on the communication-through-coherence picture [7], the window is naturally tied to an oscillation period. Mutual influence is available only during phases when sender and receiver are jointly excitable. For another example, one could even choose to represent theories of consciousness which depend on quantum entanglement [34] by setting θ to 0, although this paper does not venture down that particular path. On this reading, θ is an empirical property of a substrate and a task, not something we can dial up without changing the underlying mechanism.

Operationally, this is why the falsifiers in this paper treat θ as an input to be measured. If an advocate of a particular unity marker wants a large θ , they owe a corresponding mechanistic story and an operational measurement procedure that actually returns that large value.

The Chord postulate also imposes a necessary lower bound on any candidate window. Rearranging the diameter bound gives

$$\theta \geq \theta_{\min}^{\text{ex}} := \frac{D}{\kappa v}.$$

This is the least time budget compatible with the required within-window exchange for a support of diameter D in a substrate with signal ceiling v and exchange architecture κ . A system does not get to claim a smaller window than its own geometry and signalling allow.

There is also an upper side, provided by ingredient persistence when persistence is the mechanism invoked to close the Temporal Gap. Supplementary Note 7 proves that if every ingredient of a statement ℓ stays true once it appears within the relevant horizon, then ingredient-wise occurrence collapses into co-instantiation. This does not mean Chord logically requires persistence. Chord only requires that co-instantiation happens at least once inside the window. But persistence supplies a stability ceiling for a fixed ingredient set. If the same fine-grained grounded ingredients are supposed to belong to one moment, then the window cannot outrun the timescale over which those ingredients remain continuously well-defined. Let $\tau(\ell)$ denote the shortest persistence timescale among the grounded ingredients of ℓ . Then any persistence-supported candidate window must satisfy

$$\theta \leq \tau(\ell).$$

Put together, a persistence-supported Chord moment is only possible when there exists a nonempty feasibility interval

$$\frac{D}{\kappa v} \leq \theta \leq \tau(\ell).$$

This does *not* derive one universal exact value of θ . It derives a physically admissible range. The actual realised window is substrate- and content-relative. Primate callosal data constrain the lower side for whole-brain candidates. Ingredient lifetimes constrain the upper side whenever persistence is the mechanism that keeps the ingredients available as one moment.

This is the missing step if one worries that θ was previously a free knob. The theory does not say “pick a window you like”. It says “show that your candidate window fits inside the physical interval your own ingredients and architecture permit”. If the interval is empty, the moment is impossible.

Implications for Populations and Human-AI Hybrids

Hybrid cognition comes in progressive categories [20]. At one extreme, an AI is a prosthesis. It extends human agency but does not participate as a grounded ingredient in the human moment. This prosthetic category fails when the system becomes too much for a human to directly control. Further along, humans, AIs and prosthetic human-AI hybrids can all be distinct agents that coordinate across time. This yields a cooperative hybrid, which fails when the goals of the AI and humans involved drift out of alignment [21]. Finally there can be integrated hybrids, where human and AI contributions are meant to participate in one co-instantiated moment, extending the conscious agency of a human further into a system of AIs. An integrated hybrid can facilitate larger prosthetic and cooperative hybrids, because human agency and alignment can extend further [20].

The spacetime bound constrains this last, strongest notion of integration. If the hybrid must co-instantiate grounded ingredients and complete within-window causal exchange, then the full hybrid support must fit inside the same latency budget. If any critical part of the control loop sits behind a high latency channel, the system cannot form one unified moment at that scale. What you get instead is two unified systems taking turns across moments, not one enlarged mind.

In this sense, integrated hybrids have an integration radius. Beyond that radius you do not get a bigger mind. You get a bigger committee.

I now consider five concrete cases. The first four ask what the Chord postulate implies, so the bound is relevant. The fifth asks what remains if one drops Chord for Arpeggio, where the bound is no longer the main discriminant.

The box makes the division of labour explicit. Stack Theory says what kind of self structure must exist for consciousness of order one, two, or three. Chord then asks whether those ingredients can actually become one physically unified moment. That is why the self hierarchy and the spacetime bound complement one another rather than competing.

Case 1: Ant colonies and liquid brains. Ant colonies are a canonical example of a liquid brain [19]. Individual ants communicate through pheromone trails and physical contact. There is no persistent high-bandwidth wiring. Signal propagation through the colony relies on movement, and the effective speed is bounded by the walking speed of individual ants, on the order of centimetres per second. Colony diameters can range from metres to tens of metres.

Under Chord, a colony-scale moment would require all grounded ingredients to be co-instantiated and causally exchanged within a window θ . If the ingredients are neural-scale properties of individual ants, then θ is on the order of milliseconds, v is on the order of 10^{-2} m s^{-1} , and a colony diameter of 10 m gives $D/(v\theta) \sim 10^3$, far above any plausible κ . Equivalently, the exchange lower bound is on the order of $10^3/\kappa \text{ s}$. On that reading, the feasible interval is empty by many orders of magnitude.

One can try to rescue the colony by moving to slower, coarser ingredients such as pheromone gradients, task allocation ratios, or other colony-scale statistical properties. Those may persist for hours, which can reopen the diameter side of the feasibility interval. But that does not settle the matter. Colony-scale statistical properties are not grounded in a single microstate in the way that a neural firing pattern is. Whether they can be co-instantiated depends on the grounding resolution. Chord also still requires adequate concurrency at the chosen scale. Those requirements are not automatically supplied by long persistence. They shift the burden onto grounding, concurrency, and self-structure. So the right conclusion is stronger than mere agnosticism but weaker than a blanket impossibility claim. Ant-colony consciousness is ruled out at fine, fast grounding scales and remains heavily burdened even at slow collective scales.

Under Arpeggio, it suffices that each ingredient occurs somewhere during the window. Because ants do eventually relay information across the colony, a sufficiently long window would satisfy ingredient-wise occurrence. Arpeggio therefore permits colony-scale consciousness in principle. Chord makes that permissiveness much harder to sustain.

This does not mean individual ants cannot be conscious. An individual ant has a solid brain with persistent wiring. The bound applies to the spatial scale at which consciousness is claimed. It rules out a single conscious moment spanning the entire colony, not consciousness at the scale of individual members.

Case 2: A brain-computer interface hybrid. Consider a human with a high-bandwidth intracortical BCI, where an electrode array in motor cortex communicates bidirectionally with an external processor. Suppose the link operates at an effective round trip latency of 10 ms. The external processor sits within a few metres.

The relevant v for the link is the effective signal speed through the BCI channel. For a wired connection at near-light speed the propagation delay over a few metres is negligible, so the bottleneck is the processing and transduction latency at each end. If the round trip latency is 10 ms and θ is in

Box 2: Stack Theory of Consciousness under Arpeggio and under Chord.

Fix an organism $\alpha = \langle v_\alpha, \mu_\alpha, p_\alpha, <_\alpha \rangle$, a candidate moment statement $\ell \in L_{v_\alpha}$, and a window $\sigma = [t, t + \theta]$. The self hierarchy comes from Stack Theory and *Why Is Anything Conscious?* [10,11]. The unity test comes from the Chord and Arpeggio formalism of this paper [12]. Here I_1 says viability depends on self versus non-self discrimination. I_2 says viability depends on audience-conditioned report, so some signal must decode differently for different audience states. I_3 says viability depends on credible commitment in a trust niche with public self-binding. Write $S_n(\alpha) := \{\alpha^1, \dots, \alpha^n\}$. Let c_{conc} be the concurrency capacity at the chosen grounding scale, and let G_{ℓ_n} be the exchange graph on the grounded support of ℓ_n .

Arpeggio. Necessary conditions only.

For order 1, a conscious candidate under Arpeggio must satisfy

$$\exists \ell_1 \in L_{v_\alpha} \left(\alpha^1 \in \ell_1 \wedge \mathcal{C}(INT_\alpha, OBS_\alpha) \neq \emptyset \wedge I_1 \wedge \text{Occur}(\ell_1, \sigma) \right).$$

For order 2, it must satisfy

$$\exists \ell_2 \in L_{v_\alpha} \left(\{\alpha^1, \alpha^2\} \subseteq \ell_2 \wedge \mathcal{C}(INT_\alpha, OBS_\alpha) \neq \emptyset \wedge \exists b, \alpha^2 = c_a^{ba} \in L_{v_\alpha} \wedge I_2 \wedge \text{Occur}(\ell_2, \sigma) \right).$$

For order 3, it must satisfy

$$\exists \ell_3 \in L_{v_\alpha} \left(\{\alpha^1, \alpha^2, \alpha^3\} \subseteq \ell_3 \wedge \mathcal{C}(INT_\alpha, OBS_\alpha) \neq \emptyset \wedge \exists b, \alpha^2 = c_a^{ba} \in L_{v_\alpha} \wedge \exists b, \alpha^3 = c_a^{baba} \in L_{v_\alpha} \wedge I_3 \wedge \text{Occur}(\ell_3, \sigma) \right).$$

In plain English, the right Stack Theory self structure must exist at the chosen scale, and the relevant ingredients must each occur somewhere in the window. This is not sufficient for one unified moment, because $\text{TempGap}(\ell_n, \sigma)$ can still hold.

Chord. Sufficient conditions in this paper.

For order 1, a unified conscious moment under Chord is obtained if there exists $\ell_1 \in L_{v_\alpha}$ such that

$$\alpha^1 \in \ell_1 \wedge \mathcal{C}(INT_\alpha, OBS_\alpha) \neq \emptyset \wedge I_1 \wedge \text{CoInst}(\ell_1, \sigma),$$

with

$$c_{\text{conc}} \geq |\ell_1|, \quad D(\ell_1) \leq \kappa(G_{\ell_1})v\theta,$$

$$\eta_{u \rightarrow w} + \eta_{w \rightarrow u} \leq \theta \quad \text{for all } (u, w) \in E(G_{\ell_1}).$$

For order 2, replace the self condition by

$$\{\alpha^1, \alpha^2\} \subseteq \ell_2 \wedge \mathcal{C}(INT_\alpha, OBS_\alpha) \neq \emptyset$$

$$\wedge \exists b, \alpha^2 = c_a^{ba} \in L_{v_\alpha} \wedge I_2 \wedge \text{CoInst}(\ell_2, \sigma),$$

with the same concurrency and exchange conditions on ℓ_2 .

For order 3, replace the self condition by

$$\{\alpha^1, \alpha^2, \alpha^3\} \subseteq \ell_3 \wedge \mathcal{C}(INT_\alpha, OBS_\alpha) \neq \emptyset$$

$$\wedge \exists b, \alpha^2 = c_a^{ba} \in L_{v_\alpha} \wedge \exists b, \alpha^3 = c_a^{baba} \in L_{v_\alpha}$$

$$\wedge I_3 \wedge \text{CoInst}(\ell_3, \sigma),$$

again with the same concurrency and exchange conditions on ℓ_3 .

If persistence is the mechanism used to keep the same ingredients available, also require

$$\frac{D(\ell_n)}{\kappa(G_{\ell_n})v} \leq \theta \leq \tau(\ell_n).$$

In plain English, the same ingredients must be jointly true in one objective moment, the architecture must be able to host them together, and the required two-way exchange must fit inside the same window. Every clause is falsifiable at the chosen grounding scale.

the range of 20–50 ms, there is room for one round trip within the window. Under hub exchange this satisfies the bound with a modest margin.

This case is marginal. Chord permits integrated consciousness across the BCI link, but only if the link latency stays within the window budget. If the external processor is remote and adds network latency, pushing the round trip above θ , then the bound is violated and the system fragments into two agents that take turns. Under Arpeggio, there is no such latency ceiling. Any delay is acceptable as long as the ingredients eventually occur.

Case 3: Cloud-hosted AI. Consider a language model running on a data centre. The computation is distributed across many accelerators connected by high-speed interconnects. Typical inter-node round trip latencies in a data centre are on the order of microseconds. The physical diameter of the cluster is on the order of tens of metres. With v effectively near the speed of light in fibre and D on the order of 10^2 m, the ratio $D/(v\theta)$ is extremely small for any biologically plausible θ .

However, the Chord postulate requires co-instantiation of grounded ingredients, not just fast communication. In a conventional architecture, the system computes sequentially or in pipelined stages. At any given clock cycle, only a fraction of the model's parameters and activations are being updated. The effective concurrency capacity at the physically grounded level may be far below what is needed for all ingredients of a candidate moment to be jointly true at one instant.

This connects to the concurrency capacity result. If the architecture serialises updates so that at most c contributors are active at any clock cycle, then a moment requiring $n > c$ simultaneous contributors cannot be co-instantiated (Supplementary Note 8). The latency budget is satisfied, but co-instantiation fails. Under Chord, the system does not host a unified moment of the relevant kind. Under Arpeggio, it could, because each contributor eventually activates within the window.

This case illustrates that the Chord postulate constrains consciousness through two mechanisms. The diameter bound constrains spatially extended systems. The co-instantiation requirement constrains temporally serialised systems. Both must be satisfied.

Case 4: A population of humans. Consider a group of humans communicating via speech. Speech has an effective bandwidth much lower than neural signalling and a propagation speed limited by sound. Even ignoring speed of sound (which imposes negligible delay at the scale of a room), the bottleneck is the latency of language production and comprehension. A conversational turn takes hundreds of milliseconds to seconds.

Under Chord, any candidate integration window θ would need to accommodate a round trip of language-mediated exchange. If θ is on the order of tens of milliseconds, which is the timescale consistent with individual neural integration, then language is far too slow. Even with a generous second-scale window, the effective v of language-mediated exchange is low and the group diameter large enough that $D/(v\theta)$ exceeds the bound for groups larger than a few individuals in close proximity.

A population of humans is therefore not a single conscious entity under Chord. Each human has their own conscious moments. The population coordinates across moments, much like the cooperative hybrid category. Under Arpeggio, one could in principle argue for population-scale consciousness, since each human's contribution eventually occurs within a long enough window. Whether this is plausible is a separate question, but Arpeggio does not rule it out.

Case 5: What is conscious under Arpeggio? Arpeggio requires only that each ingredient occurs at least once somewhere inside the integration window. It imposes no constraint on co-instantiation, causal exchange, diameter, or concurrency. Without further constraints on ingredient choice and on what sets θ , nearly any physical system can be made to satisfy Arpeggio for some reading of what the ingredients are and over some timescale. A river, a weather system, a galaxy, a sufficiently old rock might be conscious. Over a long enough window, each candidate ingredient will be instantiated somewhere at some time. In the absence of further constraints, Arpeggio collapses into panpsychism [35]. Panpsychism is a coherent philosophical position, but it does not by itself distinguish conscious from non-conscious systems, which limits its use as a scientific criterion.

Stack Theory does impose further constraints (see Box 2). A conscious system must support at least a first-order self, a causal identity that separates self-generated interventions from matched observations [10,11]. This minimal, pre-reflective form of self-awareness [36] is a structural requirement on the vocabulary and the trajectory, not a temporal windowing condition. A river and a rock do not generate interventions distinguishable from the rest of their dynamics. Above the first-order self, Stack Theory derives higher orders of self (second-order for audience modelling, third-order for narrative identity) from the same weakness-maximisation principle and Bennett's Razor [37,38]. Supplementary Note 9 proves that in a system whose vocabulary is unstructured, the probability that a first-order self

candidate exists decreases exponentially in the number of distinct visited states. Only systems shaped by selection or design to maintain the right vocabulary structure can support a self and therefore be conscious on this account. This aligns with the view that conscious experiences are intrinsically linked to embodied experiencing subjects concerned with self-preservation [39].

Another contrastive case worth exploring is the Concious Turing Machine (CTM) proposed by Blum and Blum [40,41], which integrates global workspace theory with artificial intelligence. The CTM requires broadcast of information for experience. In the case of Chord, that means the CTM requires a new sort of hardware to facilitate synchronous broadcast. In the case of Arpeggio, it does not because broadcast is a matter of *what* is computed, not when. Either way the CTM still requires broadcast, but that requirement is much easier to satisfy under Arpeggio.

These cases share one logic. For the first four, estimate D , v , and a plausible θ at the grounding resolution where the ingredients are meant to live. Then ask two questions. Can the ingredients be co-instantiated at that grounding resolution? Can the latency budget $D \leq \kappa v \theta$ be met? If either answer is no, the candidate unified moment fragments. Case 5 shows the opposite extreme. Once those constraints are relaxed, Arpeggio becomes so permissive that it approaches panpsychism.

Persistence, Robustness, and Substrate Tradeoffs

Once θ is treated as a feasibility interval rather than a free constant, persistence becomes one of the main physical discriminators. It is not the only one. Grounding, concurrency, and self-structure still matter. But persistence now does real explanatory work without pretending to fix the whole problem by itself.

Supplementary Note 7 proves that if every ingredient of a statement ℓ is persistent within a window of horizon Δ , meaning that once an ingredient becomes true it stays true until the window ends, then ingredient-wise occurrence automatically implies co-instantiation. Persistence therefore closes the Temporal Gap only for windows that fit inside the relevant persistence horizon. This is why the conservative corollary

$$D \leq \kappa v \tau(\ell)$$

is best read as a ceiling on *persistence-supported* Chord moments, not as a derivation of one universal exact window.

Persistence and robustness. Different substrates have very different persistence timescales.

Consider a human brain. The grounded ingredients for a conscious moment include neural firing patterns, membrane potential configurations, and transient synaptic states. A cortical state configuration typically persists on the order of 10–100 ms before ongoing dynamics change it [15,16]. This gives $\tau(\ell)$ on the order of tens of milliseconds. With $v \approx 10 \text{ m s}^{-1}$ for myelinated callosal conduction, the derived bound gives $D \leq \kappa \cdot 0.5 \text{ m}$. Under hub exchange ($\kappa = 1$), this accommodates a human brain diameter of roughly 0.15 m with room to spare.

Now consider damage. Delete a random half of a human brain and most grounded ingredients are destroyed outright. The surviving tissue may support a conscious moment of reduced content, but the original moment cannot be reconstituted. A solid brain is brittle. Its persistence timescale is fragile with respect to damage.

An ant colony is the opposite. The grounded ingredients at the colony scale, if they exist, would be statistical properties of the collective such as pheromone gradients, foraging trail distributions, and task allocation ratios. Delete a random half of the colony with your foot and most of these properties survive, because they are distributed across many redundant agents and are actively repaired by the survivors. The persistence timescale $\tau(\ell)$ for colony-scale ingredients is on the order of hours to days and is robust to large perturbations. This robustness is a form of self-repair, which connects to work on why living systems exist [11,19,42–44]. Self-repair is the mechanism by which $\tau(\ell)$ becomes long enough to support extended integration windows.

Worked examples under the derived bound.

1. **Human brain (fine-grained neural ingredients).** $v \approx 10 \text{ m s}^{-1}$, $\tau(\ell) \approx 50 \text{ ms}$. Derived diagnostic: $D \leq \kappa \cdot 0.5 \text{ m}$. Human brain diameter $\approx 0.15 \text{ m}$. Compatible with Chord under both hub exchange ($\kappa = 1$) and all-to-all ($\kappa = \frac{1}{2}$). Consistent with consciousness at the whole-brain scale.
2. **Ant colony (coarse-grained colony-scale ingredients).** $v \approx 10^{-2} \text{ m s}^{-1}$, $\tau(\ell) \approx 10^4 \text{ s}$ (hours, for colony-scale statistical properties). Derived diagnostic: $D \leq \kappa \cdot 10^2 \text{ m}$. A small colony ($D \approx 10 \text{ m}$) can satisfy the diameter budget under hub exchange. Whether Chord is satisfied still depends on co-instantiation and concurrency at the relevant grounding resolution, not just on diameter.

The derived bound $D \leq \kappa v \tau(\ell)$ reveals a tradeoff between speed and persistence. Solid brains have high v but short, fragile $\tau(\ell)$. Liquid brains have low v but potentially long, robust $\tau(\ell)$. The product $v \tau(\ell)$ helps determine the spatial reach of a persistence-supported unified moment, but it does not settle the whole story by itself.

The moral is not that persistence alone decides consciousness. Rather, θ , v , κ , grounding resolution, and concurrency capacity jointly constrain what kinds of unified moments are even physically available. Persistence matters because “just make θ bigger” is not cost-free. It typically forces a move to coarser, slower ingredients and thereby changes the hypothesised phenomenology.

Under Arpeggio, that persistence advantage becomes even more consequential. If no co-instantiation is needed, then slow but persistent systems like ant colonies or oceanic current systems become candidate conscious systems, operating on timescales from hours to geological epochs. Chord is more selective. Once the window is no longer free, $\tau(\ell)$ becomes one of the main discriminative physical quantities, but not the only one. Grounding, concurrency, and self-structure still matter.

Predictions and Falsifiers

Each falsifier targets a different modelling commitment.

1. **Temporal Gap falsifier.** Take any high time resolution recording where you can define a set of grounded ingredients ℓ and a candidate unity marker over windows of duration θ . A unity marker is any measurable variable that is supposed to indicate that one unified moment occurred, such as a phase synchrony measure or a behavioural report. For each window $[t, t + \theta]$, compute $\text{Occur}(\ell, [t, t + \theta])$ and $\text{CoInst}(\ell, [t, t + \theta])$ for the same ℓ . Then focus on windows where Occur holds but CoInst fails. If the unity marker still behaves as if a single moment occurred in those windows, then the Chord requirement is wrong for that marker.
2. **Latency budget violation.** Pick any substrate where you can estimate a support diameter and a propagation ceiling for the same candidate unity marker. Write D_L for a conservative lower bound on diameter, and write v_U and θ_U for conservative upper bounds on speed and window duration. Compute the conservative margin $M_{\text{cons}} = D_L - \kappa v_U \theta_U$. If M_{cons} is significantly positive across windows under the protocol in Supplementary Note 5, then either the within-window exchange postulate is false or the marker is not a unified moment.
3. **Architecture factor shift.** Hold the same nodes, the same geometry, and the same window definition. Change only the exchange graph. For engineered systems this means swapping between hub exchange, all-to-all exchange, and a chosen sparse graph. For the same measured v and θ , the largest diameter that still supports integration should scale in direct proportion to κ . If the transition does not move when κ moves, then the architecture factor model is wrong.

These are just sense checks on timestamps, distances, and who must talk to whom. A unified moment is the set of sites that can mutually exchange causal influence within a window of duration θ . Its spatial diameter cannot exceed $\kappa v \theta$. If you want a bigger subjective moment, you must either have more time for integration (somehow), propagate faster, or change the exchange architecture.

Methods

Stack Theory Objects and Temporal Lift

Definitions 1 and 2 give the formal objects used in the Results. I repeat the core pieces here for convenience.

An environment is a nonempty set Φ of mutually exclusive microstates. A program is a set $p \subseteq \Phi$. A finite vocabulary $\mathfrak{v} \subseteq 2^\Phi$ is the set of programs a system can implement. The induced embodied language is

$$L_{\mathfrak{v}} := \{ \ell \subseteq \mathfrak{v} \mid \bigcap_{p \in \ell} p \neq \emptyset \}.$$

For $\ell \in L_{\mathfrak{v}}$, the truth set is $T(\ell) := \bigcap_{p \in \ell} p$, with $T(\emptyset) = \Phi$.

To lift truth to time, let $\phi(\tau) \in \Phi$ be the microstate trajectory. Arpeggio, Chord, and the related synchrony conditions are defined in Definitions 2 and 3. Supplementary Note 7 gives the exact Appendix 26 discrete-time lift. There an objective trajectory $\phi : \mathbb{N} \rightarrow \Phi$ induces windows $\sigma_u = (\phi(u), \dots, \phi(u+\Delta)) \in \Phi_{\Delta}$, with $\text{Occur}(\ell, \sigma_u) \iff \sigma_u \in T(\diamond_{\Delta}\ell)$ and $\text{CoInst}(\ell, \sigma_u) \iff \sigma_u \in \diamond_{\Delta}T(\ell)$. The interval notation in the main text is just the continuous-time reading of the same distinction. Grounded means that the programs are evaluated directly on $\phi(\tau)$ in the base environment.

Architecture Factor and Proof Idea

Supplementary Note 3 defines $\kappa = \frac{1}{2}h(G)$ where $h(G)$ is the hop diameter of an exchange graph G on the grounded support. In other words, $h(G)$ counts the worst case number of message handoffs needed to connect the two most separated sites, given the required exchange edges.

Theorem 1 follows from one inequality. Assume each edge (u, w) can complete a round trip within the window so $\eta_{u \rightarrow w} + \eta_{w \rightarrow u} \leq \theta$. Then at least one direction has travel time at most $\theta/2$. With propagation speed bounded by v this implies the metric distance between u and w is at most $v\theta/2$.

Any pair of sites is connected by a path of at most $h(G)$ hops. The triangle inequality says the direct distance between two sites is at most the sum of distances along any path between them. So their separation is at most $h(G)v\theta/2 = \kappa v\theta$.

Mechanistic Integration Model

The mechanistic model places N contributor sites uniformly on a circle of diameter D . For all-to-all exchange I require a completed round trip for every unordered pair of sites, using the Euclidean chord distance on the circle. Chord distance is the straight line distance between two points on the circle. For hub exchange I model a central mediator at the circle centre, so each contributor must complete a round trip via the hub.

A directed message time is $d/v + \zeta$, where d is the one-way Euclidean distance and ζ is exponential jitter with mean $j\theta$. A round trip time is therefore $2d/v + \zeta_1 + \zeta_2$. Under the independence assumptions, overall success probability factorises over required exchanges.

Because $\zeta_1 + \zeta_2$ is the sum of two independent exponential random variables, it has a Gamma distribution. I therefore evaluate success probability exactly using the closed form Gamma cumulative distribution function. In other words, this is just the probability that the random delay stays below the remaining time budget.

Primate Literature Extraction

The primate dataset is extracted from Table 1 of Phillips et al [25] plus the published correction [26]. The corrected Table 1 contains $n = 15$ individuals across 14 species. I include the extracted values in `ttgs_phillips2015_table1.csv`. I treat the reported interhemispheric conduction times as θ_{\min} proxies for moments requiring bilateral exchange.

Monte Carlo Stress Tests

I simulate a measurement pipeline where D , θ , and v are observed with noise. For each simulated run I form the margin $M = \widehat{D} - \kappa \widehat{v} \widehat{\theta}$ across windows and test whether its mean is positive with a one sided t test. In other words, this asks whether the average margin is reliably above zero given sampling noise.

I treat v as a sampled quantity and compare estimators. I also simulate dependence across windows with an AR(1) model. AR(1) means a first order autoregressive process, where each window is correlated with the previous one. I show that thinning controls false refutations. Thinning means keeping only every k th window so that the remaining windows are closer to independent. For example, if you slide a 50 ms window forward in 1 ms steps, then adjacent windows share almost all of their data. Keeping every 50th window removes that overlap.

Author Contributions: Michael Timothy Bennett: Conceptualization, Formal analysis, Investigation, Methodology, Software, Visualization, Writing – original draft, Writing – review & editing.

Funding: No external funding was received for this work.

Data Availability Statement: All simulation outputs and source data tables are generated by `ttgs_simulation.py` and written into the `ttgs_source_data` directory. The primate dataset extracted from published literature is included as `ttgs_phillips2015_table1.csv`. Source data are provided with this paper.

Acknowledgments: The author thanks colleagues and early readers for feedback that improved the manuscript.

Conflicts of Interest: The author declares no competing interests.

Code Availability: All code required to reproduce the figures and tables is provided in `ttgs_simulation.py`. Running `python ttgs_simulation.py` regenerates every figure, LaTeX macro file, and source data CSV used by this manuscript. The same run writes `ttgs_provenance.json`, a machine readable record of every literature derived numeric input and its stable identifier. A one-page reproducibility and data provenance checklist is included with the submission package.

References

1. Ned Block. On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18(2):227–247, 1995. doi:10.1017/s0140525x00038188.
2. Bernard J. Baars. *A Cognitive Theory of Consciousness*. Cambridge University Press, Cambridge, 1988.
3. Stanislas Dehaene and Lionel Naccache. Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, 79(1-2):1–37, 2001. doi:10.1016/s0010-0277(00)00123-2.
4. Victor Lamme. Towards a true neural stance on consciousness. *Trends in cognitive sciences*, 10(11):494–501, 2006. doi:10.1016/j.tics.2006.09.001.
5. Francisco J. Varela, Jean-Philippe Lachaux, Eugenio Rodriguez, and Jacques Martinerie. The brainweb: phase synchronization and large-scale integration. *Nature Reviews Neuroscience*, 2(4):229–239, 2001. doi:10.1038/35067550.
6. Wolf Singer and Charles M. Gray. Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience*, 18:555–586, 1995. doi:10.1146/annurev.ne.18.030195.003011.
7. Pascal Fries. A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, 9(10):474–480, 2005. doi:10.1016/j.tics.2005.08.011.
8. Giulio Tononi. An information integration theory of consciousness. *BMC Neuroscience*, 5(1):42, 2004. doi:10.1186/1471-2202-5-42.
9. Michael Timothy Bennett. Emergent causality and the foundation of consciousness. In *16th International Conference on Artificial General Intelligence*, Lecture Notes in Computer Science, pages 52–61. Springer, 2023. OUCI metadata page: <https://ouci.dntb.gov.ua/en/works/7BoXJMW4/>. doi:10.1007/978-3-031-33469-6_6.
10. Michael Timothy Bennett, Sean Welsh, and Anna Ciaunica. *Why Is Anything Conscious?* Preprint, 2024. URL: <https://arxiv.org/abs/2409.14545>, arXiv:2409.14545, doi:10.48550/arxiv.2409.14545.
11. Michael Timothy Bennett. *How To Build Conscious Machines*. PhD thesis, The Australian National University, 2025. URL: <https://hdl.handle.net/1885/733782452>, doi:10.25911/ta58-p428.

12. Michael Timothy Bennett. A mind cannot be smeared across time. In *Forthcoming in Proceedings of the AAAI 2026 Spring Symposium on Machine Consciousness: Integrating Theory, Technology, and Philosophy*, Burlingame, California, USA, April 2026. AAAI Press.
13. Daniel C. Dennett and Marcel Kinsbourne. Time and the observer: the where and when of consciousness in the brain. *Behavioral and Brain Sciences*, 15(2):183–201, 1992. doi:10.1017/s0140525x00068229.
14. Anna Ciaunica, Andreas Roepstorff, Aikaterini Katerina Fotopoulou, and Bruna Petreca. Whatever next and close to my self—the transparent senses and the “second skin”: Implications for the case of depersonalization. *Frontiers in Psychology*, 12, 2021. doi:10.3389/fpsyg.2021.613587.
15. David M. Eagleman and Terrence J. Sejnowski. Motion integration and postdiction in visual awareness. *Science*, 287(5460):2036–2038, 2000. doi:10.1126/science.287.5460.2036.
16. Ernst Pöppel. A hierarchical model of temporal perception. *Trends in Cognitive Sciences*, 1(2):56–61, 1997. doi:10.1016/s1364-6613(97)01008-5.
17. Jean Vroomen and Mirjam Keetels. Perception of intersensory synchrony: a tutorial review. *Attention, Perception, & Psychophysics*, 72(4):871–884, 2010. doi:10.3758/app.72.4.871.
18. Mark T. Wallace and Ryan A. Stevenson. The construct of the multisensory temporal binding window and its dysregulation in developmental disabilities. *Neuropsychologia*, 64:105–123, 2014. doi:10.1016/j.neuropsychologia.2014.08.005.
19. Ricard Solé, Melanie Moses, and Stephanie Forrest. Liquid brains, solid brains. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 374(1774):20190040, 2019. URL: <https://royalsocietypublishing.org/doi/abs/10.1098/rstb.2019.0040>, arXiv:<https://royalsocietypublishing.org/doi/pdf/10.1098/rstb.2019.0040>, doi:10.1098/rstb.2019.0040.
20. Ricard Solé, Luis F. Seoane, Jordi Pla-Mauri, Michael Timothy Bennett, Michael E. Hochberg, and Michael Levin. Cognition spaces: natural, artificial, and hybrid, 2026. URL: <https://arxiv.org/abs/2601.12837>, arXiv:2601.12837, doi:10.48550/arxiv.2601.12837.
21. Michael Timothy Bennett. Are biological systems more intelligent than artificial intelligence?, 2024. In press, 2026, *Philosophical Transactions of the Royal Society B: Biological Sciences*. Special issue on Hybrid agencies: crossing borders between biological and artificial worlds. URL: <https://arxiv.org/abs/2405.02325>, arXiv:2405.02325, doi:10.48550/arxiv.2405.02325.
22. Jacob D. Bekenstein. Universal upper bound on the entropy-to-energy ratio for bounded systems. *Phys. Rev. D*, 23:287–298, January 1981. URL: <https://doi.org/10.1103/PhysRevD.23.287>, doi:10.1103/physrevd.23.287.
23. Michael Timothy Bennett. Is complexity an illusion? In *17th International Conference on Artificial General Intelligence*, Lecture Notes in Computer Science. Springer, 2024. doi:10.1007/978-3-031-65572-2_2.
24. Michael Timothy Bennett. Technical appendices, 2025. Archived release on Zenodo. Source repository: <https://github.com/ViscousLemming/Technical-Appendices>. doi:10.5281/zenodo.7641741.
25. Kimberley A. Phillips, Cheryl D. Stimpson, Jeroen B. Smaers, Mary Ann Raghanti, Bob Jacobs, Aleksandar Popratiloff, Patrick R. Hof, and Chet C. Sherwood. The corpus callosum in primates: processing speed of axons and the evolution of hemispheric asymmetry. *Proceedings of the Royal Society B: Biological Sciences*, 282(1818):20151535, 2015. doi:10.1098/rspb.2015.1535.
26. Kimberley A. Phillips, Cheryl D. Stimpson, Jeroen B. Smaers, Mary Ann Raghanti, Bob Jacobs, Aleksandar Popratiloff, Patrick R. Hof, and Chet C. Sherwood. Correction to The corpus callosum in primates: processing speed of axons and the evolution of hemispheric asymmetry. *Proceedings of the Royal Society B: Biological Sciences*, 282(1819):20152620, 2015. doi:10.1098/rspb.2015.2620.
27. John L. Ringo, Richard W. Doty, Steven Demeter, and Pierre Y. Simard. Time is of the essence: a conjecture that hemispheric specialization arises from interhemispheric conduction delay. *Cerebral Cortex*, 4(4):331–343, 1994. doi:10.1093/cercor/4.4.331.
28. Roberto Caminiti, Houda Ghaziri, Ralf A. W. Galuske, Patrick R. Hof, and Giorgio M. Innocenti. Evolution amplified processing with temporally dispersed slow neuronal connectivity in primates. *Proceedings of the National Academy of Sciences of the United States of America*, 106(46):19551–19556, 2009. doi:10.1073/pnas.0907655106.
29. Giorgio M. Innocenti, Ingo Vahlsing, and Roberto Caminiti. The functional characterization of callosal connections. *Progress in Neurobiology*, 208:102186, 2022. doi:10.1016/j.pneurobio.2021.102186.
30. Lucia Melloni, Carlos Molina, Miguel Pena, David Torres, Wolf Singer, and Eugenio Rodriguez. Synchronization of neural activity across cortical areas correlates with conscious perception. *Journal of Neuroscience*, 27(11):2858–2865, 2007. doi:10.1523/jneurosci.4623-06.2007.

31. Marcello Massimini, Fabio Ferrarelli, Reto Huber, Steven K. Esser, Harminder Singh, and Giulio Tononi. Breakdown of cortical effective connectivity during sleep. *Science*, 309(5744):2228–2232, 2005. doi:10.1126/science.1117256.
32. Adenauer G. Casali, Olivia Gosseries, Mario Rosanova, M'elanie Boly, Simone Sarasso, Karina R. Casali, Silvia Casarotto, Marie-Aur'elie Bruno, Steven Laureys, Giulio Tononi, and Marcello Massimini. A theoretically based index of consciousness independent of sensory processing and behavior. *Science Translational Medicine*, 5(198):198ra105, 2013. doi:10.1126/scitranslmed.3006294.
33. Jonathon Sendall. Event horizons, spacetime geometry, and the limits of integrated consciousness, 2026. URL: <https://arxiv.org/abs/2512.23105>, arXiv: 2512. 23105.
34. Stuart Hameroff and Roger Penrose. Consciousness in the universe: A review of the 'orch or' theory. *Physics of Life Reviews*, 11(1):39–78, 2014. URL: <https://www.sciencedirect.com/science/article/pii/S1571064513001188>, doi:10.1016/j.plrev.2013.08.002.
35. Galen Strawson. Realistic monism - why physicalism entails panpsychism. *Journal of Consciousness Studies*, 13(10-11):3–31, 2006.
36. Anna Ciaunica and Laura Crucianelli. Minimal self-awareness: from within a developmental perspective. *Journal of Consciousness Studies*, 26(3-4):207–226, 2019. URL: <https://www.ingentaconnect.com/content/imp/jcs/2019/00000026/f0020003/art00010>.
37. Michael Timothy Bennett. The optimal choice of hypothesis is the weakest, not the shortest. In *16th International Conference on Artificial General Intelligence*, Lecture Notes in Computer Science, pages 42–51. Springer, 2023. doi:10.1007/978-3-031-33469-6_5.
38. Michael Timothy Bennett. Optimal policy is weakest policy. In *Artificial General Intelligence*, volume 16057 of *Lecture Notes in Computer Science*, pages 43–56. Springer, 2025. doi:10.1007/978-3-032-00686-8_5.
39. Anna Ciaunica, Adam Safron, and Jonathan Delafield-Butt. Back to square one: the bodily roots of conscious experiences in early life. *Neuroscience of Consciousness*, 2021(2), 11 2021. doi:10.1093/nc/niab037.
40. Manuel Blum and Lenore Blum. A theoretical computer science perspective on consciousness. *J. Artif. Intell. Conscious.*, 8:1–42, 2020.
41. Lenore Blum and Manuel Blum. A theory of consciousness from a theoretical computer science perspective: Insights from the conscious turing machine. *Proceedings of the National Academy of Sciences*, 119(21):e2115934119, 2022. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.2115934119>, arXiv:<https://www.pnas.org/doi/pdf/10.1073/pnas.2115934119>, doi:10.1073/pnas.2115934119.
42. Michael Timothy Bennett. Why does life exist?, 2026. Preprint, 2026. Available at <https://www.preprints.org/manuscript/202603.0203>.
43. Ricard Solé, Christopher P Kempes, Bernat Corominas-Murtra, Manlio De Domenico, Artemy Kolchinsky, Michael Lachmann, Eric Libby, Serguei Saavedra, Eric Smith, and David Wolpert. Fundamental constraints to the logic of living systems. *Interface Focus*, 14(5):20240010, October 2024. doi:10.1098/rsfs.2024.0010.
44. Ricard Solé and Luís F Seoane. Evolution of brains and computers: The roads not taken. *Entropy*, 24(5):665, 2022. doi:10.3390/e24050665.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.