The Speed of Imagination: A Multiscale Kinematics of Cortical Waves and Representational Transitions

PIATRA . INSTITUTE

July 2026

Abstract

Imagination crosses arbitrary distances in a fraction of a second. Picture the desk in front of you, then a galaxy ten billion light-years away; the second image arrives about as fast as the first and stays inside your skull. Dividing the represented distance by that moment gives a number far above the speed of light, and the number describes nothing that moved. This paper makes the artifact precise and replaces it with quantities that can be measured. A velocity is defined only once a changing state, a clock, and a metric on the states are fixed, and the phrase "speed of imagination" runs at least five spaces together: the cortical sheet, the neural population state, the represented content, the represented external world, and represented time. The maps between them do not preserve distance. Scaling the imagined universe by a factor λ multiplies the apparent velocity by λ while the neural trajectory is untouched, so the apparent velocity is a property of the labelling and not of the process. What is real is a finite-time transition in local neural dynamics with distance-independent representational reach, and it has no single speed. It has a kinematic profile: cortical wave propagation, population state-space velocity, target-representation latency, replay rate, temporal compression, and fidelity, measured in the literature from tens of milliseconds to a second and from centimetres per second to metres per second, none convertible into the others. Building on evidence that prefrontal traveling waves correspond to rotational population trajectories, the paper takes wave geometry to act as a control field that gates which low-dimensional subspace can become active, so that the latency of an imaginative transition is set by the control cost of the neural path rather than the represented distance. A minimal control model makes the claims concrete. Transition latency is invariant to represented-space rescaling across 20 orders of magnitude while the apparent velocity scales exactly with the label; latency is predicted by control energy at the characteristic timescale ($R^2 = 0.87$) and not by represented physical distance ($R^2 = 0.02$); a direct jump has a latency independent of the number of explicitly simulated intermediate states while a traversal is affine in that number; and forcing a transition faster than the control budget affords costs a fidelity governed by the efficiency of the neural path. An experiment that varies represented physical, semantic, visual, and graph distance orthogonally, contrasts a jump with a traversal, and reads out decoding onset, wave topology, and state-space trajectory turns the framework into falsifiable predictions, among them that perception and imagery run cortical waves in opposite directions. Relativity holds throughout, because a representational jump sends nothing to the place it represents and takes nothing back.

Two hundred milliseconds to a distant galaxy

You can hold the desk in front of you, and then, with no sense of travel, a galaxy ten billion light-years away. Someone determined to find a speed in this divides the represented distance by the time the switch took,

$$v_{\text{apparent}} = \frac{D_{\text{represented}}}{\tau_{\text{transition}}} \approx \frac{10^{10}\ \text{light-years}}{0.2\ \text{s}},$$

and reads off a value many orders of magnitude above c. The arithmetic is sound and it measures nothing that moved. Nothing left the skull. The brain passed from one state to another and both states occupied the same 20 centimetres of tissue. A digital analogue is exact: setting a variable from $x = 1$ to $x = 10^{100}$ takes one processor cycle, and no object crossed 10^{100} metres. The represented magnitude changed while the machine did a bounded amount of local work.

Going from the desk to the galaxy changes the reference. It moves nothing through the space referred to. A card makes the separation plain. You can picture a sealed card resting on Mars at once, and you cannot say which number is printed on it, because what you hold is a representation of the card and not a channel to it. Reference is not retrieval, and a mental pointer to a place delivers no information from it. So imagination combines two properties that ordinary motion never puts together. Its reach is nearly independent of the represented distance, and its update rate is finite and unspectacular. Reporting the ratio of the two as a speed treats a semantic interval as a physical one.

The productive response keeps the phenomenon and drops the word speed. The apparent velocity fails as a physical rate, and the intuition beneath it points at a set of quantities that are measurable and mechanistically real: how long a target image takes to form, how fast a scene can be transformed, how quickly a sequence of remembered states can be replayed, how far and how coherently activity propagates across cortex while all of this happens. This paper defines those quantities, argues that they do not collapse into one, and proposes what sets the most important of them.

A speed needs a space and a metric

A speed is not defined until a changing state, a clock, and a metric on the states are all in hand. The phrase "speed of imagination" leaves the metric unstated and slides between the states it might mean. At least five spaces are in play, each with its own natural distance.

Space	State variable	Natural distance
Cortical sheet $\mathcal{C}$	electrode or source position $\mathbf{r}$	millimetres
Neural population $\mathcal{N}$	firing-rate or field state $\mathbf{x}$	Mahalanobis or manifold distance
Representational $\mathcal{R}$	decoded content $\mathbf{z}$	neural or behavioural dissimilarity
Represented world $\mathcal{P}$	imagined location X	metres to light-years
Represented time $\mathcal{T}$	simulated duration T_s	seconds to geological epochs

The chain of maps $\mathcal{C} \to \mathcal{N} \to \mathcal{R} \to \mathcal{P}$ has no obligation to preserve distance at any step, and it does not. Two neural states a hair apart can represent places billions of light-years apart, and two nearly identical facial expressions can demand well-separated neural states. The represented distance $d_{\mathcal{P}}(A, B)$ and the neural distance $d_{\mathcal{N}}(A, B)$ therefore vary independently. Making the

neural side quantitative already requires a choice of metric, whether a cross-validated representational dissimilarity (Kriegeskorte, Mur, & Bandettini, 2008) or a learned graph distance over a relational environment (Behrens et al., 2018), and the choice is consequential rather than cosmetic.

The rescaling argument turns this into a clean impossibility. Write the apparent velocity as above, then multiply the scale of the imagined universe by an arbitrary factor,

$$D'_{\text{represented}} = \lambda\, D_{\text{represented}}.$$

The neural transition that implements the thought is unchanged, so its duration is unchanged, $\tau' = \tau$, and

$$v'_{\text{apparent}} = \lambda\, v_{\text{apparent}}.$$

A quantity that changes when the represented coordinates are relabelled, while every physical variable in the brain holds fixed, is a property of the labelling. The represented distance can be inflated without bound by imagining a larger universe, and the transition does not notice. This has an experimental form: teach two groups the same relational environment, tell one that neighbouring locations lie a metre apart and the other that they lie a million light-years apart, and the apparent external velocity differs by 18 orders of magnitude while the neural trajectories and the latencies, if the account here is right, coincide. The simulation below works this out inside a model whose latency is computed from the neural dynamics alone, where the represented-space map never enters the calculation and scaling it from 1 to 10^{20} leaves the latency bit-identical.

A kinematic profile in place of a speed

Once the represented distance is set aside, several genuine rates remain, and the literature has already measured them at scales that do not reconcile. Deliberately forming a specific visual image runs on hundreds of milliseconds. In a magnetoencephalography study where participants imagined previously viewed items, a classifier trained on perception decoded the imagined category from about 350 ms after the cue, a classifier trained within imagery reached significance from about 540 ms, and imagery decoding peaked near 1073 ms (Dijkstra et al., 2018). Transforming an image has its own rate: reaction time in mental rotation grows linearly with angular disparity, at roughly 60 degrees per second (Shepard & Metzler, 1971), a speed through a transformation space with no counterpart in physical space. Replaying a remembered sequence is faster still and measures something else again, with human sequence transitions peaking near 40 ms of state-to-state lag (Liu et al., 2019) and forward replay peaking near 30 ms (Huang et al., 2024). That same replay study, having detected 30 ms transitions, states plainly that it could not identify when each image was imagined or "the speed of imagination," which locates the gap from inside the field.

These numbers cannot be folded into one velocity. A 30 ms replay transition is not a 500 ms conscious image, and neither is a 60-degree rotation. The physical substrate underneath them is slower

than any of the mental rates suggests and far slower than light: long-range association fibres in the adult human brain transmit at 3 to 6 m/s with minimum latencies near 25 ms (Van Blooijs et al., 2023), which against the invariant $c \approx 3.0 \times 10^8$ m/s of special relativity (Einstein, 1905) is 50 to 100 million times slower. Cognition is not the conduction of a single fibre, but the comparison settles the physical question: the implementation is emphatically subluminal.

The construct is a profile with several axes. Collect the quantities that a transition from A to B actually has,

$$\text{IKP} = \{\, \tau,\ L_N,\ \bar{v}_N,\ W(t),\ C_T,\ F \,\},$$

the target latency τ, the neural path length L_N and mean state-space speed $\bar{v}_N$, the cortical wave state $W(t)$, the temporal compression C_T, and the fidelity F. Temporal compression is its own axis: representing "a million years pass" inside one second is a ratio

$$C_T = \frac{\Delta T_{\text{represented}}}{\Delta t_{\text{physical}}} \approx 10^6 \text{ years/second},$$

achieved by omitting almost every intermediate state. The brain does not simulate a million years at speed. Fidelity keeps a fast vague image from being scored as a fast detailed one, and a natural throughput measure ties the two together as the recoverable information about the intended scene per unit time, $\mathcal{R} = I(S;\hat{S})/T$ (Cover & Thomas, 2006), with $\hat{S}$ the scene decoded from neural activity. A person can "imagine a galaxy" almost at once only by invoking a compressed token; building a detailed, stable, inspectable galaxy costs more time and more of the profile.

Cortical waves as a control field

The most promising mechanism for what sets these rates comes from treating cortical oscillations as traveling waves rather than as synchronous rhythms. Write the oscillatory field as

$$u(\mathbf{r}, t) = A(\mathbf{r}, t)\, \cos \phi(\mathbf{r}, t),$$

with local wave vector $\mathbf{k} = \nabla_{\mathbf{r}}\phi$ and angular frequency $\omega = -\partial\phi/\partial t$, so that the phase velocity is

$$\mathbf{v}_\phi = \frac{\omega}{\|\mathbf{k}\|^2}\,\mathbf{k}, \qquad v_\phi = \frac{|\omega|}{\|\mathbf{k}\|}.$$

In lateral prefrontal cortex the theta, alpha, and beta bands organize as traveling waves with speeds of 20 to 60 cm/s that increase with frequency, more often rotating than planar, with a delay-period bias toward one propagation direction that is strongest in beta (Bhattacharya et al., 2022). A rotating wave has a center, a local direction that varies across the sheet, a curvature, and a handedness, so a single scalar speed discards most of what it is doing. A fuller description is a field,

$$W(t) = \{\, A,\ f,\ \mathbf{v}_\phi,\ \nabla \cdot \mathbf{v}_\phi,\ \nabla \times \mathbf{v}_\phi,\ q,\ \rho \,\},$$

with divergence separating sources from sinks, curl measuring rotation, q the topological charge, and ρ the spatial coherence. These are not abstractions imposed on the data. Human intracranial recordings resolve the same patterns, planar, spiral, and concentric source-and-sink waves along with complex ones, and their direction and strength distinguish behavioural states and in some electrode clusters individual remembered items (Das et al., 2026).

Phase velocity is not information velocity, and the paper keeps them apart by vocabulary. A phase maximum can move because activity is relayed between neighbours, because fixed phase offsets across independent oscillators produce an apparent sweep, or because recurrent interaction builds an emergent gradient, and at the scalp because source mixing manufactures propagation that is not there. That last effect is not hypothetical: modeling shows cortical source mixing can inflate the apparent traveling-wave speed measured at sensors while leaving the inferred direction relatively stable (Orczyk & Kajikawa, 2021), and the general point that a moving phase is not a moving signal is standard in the traveling-wave literature (Muller, Chavane, Reynolds, & Sejnowski, 2018). Cortical phase velocity, wavefront velocity, population state velocity, and content-decoding latency are four quantities, and only the last two speak to when an image is available.

What the waves do, on the evidence, is control content rather than carry it. In the prefrontal study, task state altered wave direction while the specific remembered item did not reliably change the waves, which suggests a content-independent function (Bhattacharya et al., 2022). The spatial-computing account makes this concrete: population activity lives in low-dimensional subspaces, alpha and beta impose spatially structured inhibition that determines where gamma and spiking can be expressed, and traveling waves help switch between one subspace organization and another (Miller, Brincat, & Roy, 2024). A wave here is closer to a moving stencil than to a message in a cable, opening the region where a particular computation is allowed. For imagination the mechanism reads as a sequence: a memory cue drives retrieval, wave reconfiguration opens the target subspace, sensory cortex reinstates the stored pattern top-down (Kosslyn, Ganis, & Thompson, 2001), and a stable image results.

The strongest current evidence ties wave motion to state-space motion. In macaque prefrontal cortex recovering from a working-memory distractor, population spiking follows rotational trajectories in state space, traveling waves of spiking cross the cortical surface, and the two are aspects of one dynamics: removing the leading rotational components from the population activity disrupts the wave, raising the angular deviation of the wavefront by 31.32 with a temporal consistency of 0.53, and trials whose trajectories complete a fuller rotation are the ones performed correctly (Batabyal et al., 2026). Spatial motion across tissue and abstract motion through population state space are, at least in this regime, two readings of the same event. The link from wave to neural state has partial support. The link from neural state to imagined content, its latency and its fidelity, is the piece that is missing, and it is what would make the project mechanistic rather than terminological.

Latency is a control problem

Model the low-dimensional neural state $\mathbf{z}(t) = E[\mathbf{x}(t)]$ as a controlled dynamical system,

$$\dot{\mathbf{z}} = f(\mathbf{z}) + B(\mathbf{z})\,\mathbf{u}_W(t) + \epsilon(t),$$

with f the intrinsic recurrent dynamics, $\mathbf{u}_W$ the wave-mediated control, B the map from wave states to the manifold, and ϵ noise. Reaching an imagined content means steering the state from an initial basin $\mathcal{B}_A$ into a target basin $\mathcal{B}_B$. Two quantities characterize the transition: the minimum time under bounded control,

$$T^*(A,B) = \min_{\mathbf{u}_W} \{\, T : \mathbf{z}(T) \in \mathcal{B}_B,\ \|\mathbf{u}_W(t)\| \le U_{\max} \,\},$$

and the minimum control energy, $E^*(A,B) = \min_{\mathbf{u}_W} \int_0^T \|\mathbf{u}_W(t)\|^2\, dt$, which for a linear system is read off the controllability Gramian and has been used at network scale to quantify how hard one brain state is to reach from another (Gu et al., 2015). The neural path length and mean speed follow from a metric G on the manifold,

$$v_N(t) = \sqrt{\dot{\mathbf{z}}^\top G(\mathbf{z})\,\dot{\mathbf{z}}}, \qquad L_N = \int_{t_0}^{t_1} v_N(t)\, dt,$$

where G must be a whitened or learned metric rather than raw Euclidean distance in principal-component coordinates, since otherwise rescaling a dimension rescales the apparent speed. The population geometry that makes this well posed, low-dimensional manifolds carrying rotational trajectories, is by now a standard description of cortical activity (Churchland et al., 2012; Gallego, Perich, Miller, & Solla, 2017).

Operationally the latency is when a decoder for the target crosses and holds a threshold,

$$\tau_{A\to B} = \inf \{\, t : P(B \mid \mathbf{x}(t)) > \theta \text{ for a duration } \delta \,\} - t_{\text{cue}},$$

with θ preregistered and the onset better estimated by a change-point model than by the first significant sample. The central hypothesis is that this latency is a function of the neural transition and not of the represented distance,

$$\tau_{A\to B} \approx h(L_N,\ E^*,\ \text{familiarity},\ \text{complexity},\ \text{required fidelity}), \qquad \frac{\partial \tau}{\partial D_{\text{physical represented}}} \approx 0$$

in a direct-jump condition. The claim is falsifiable, and its falsifier is a measured dependence of jump latency on represented distance once familiarity and complexity are controlled.

Imagination is not one operation, and the model expects at least three policies. A jump reaches the target by associative pattern completion, with no obligation to represent intervening positions, and its cost is a single transition. A traversal makes the intermediate states part of the task, so its duration grows with their number, roughly $\tau_{\text{traverse}} = \tau_0 + n\delta$ for n represented states at mean transition time δ, the regime where the tens-of-milliseconds replay lags apply. A transformation follows a continuous trajectory through a constrained feature manifold, where mental rotation lives and a transformation rate exists once the metric is fixed. The distinction between jumping and traversing is the single most consequential behavioural split in the framework, because it is what lets one mind reach a distant galaxy at once, take seconds to imagine travelling there, and longer still to build the world in detail. These are different control policies over the same manifold. They are not competing measurements of one speed.

What a minimal model shows

To make the claims concrete and reproducible, the accompanying simulation instantiates the control model as a linear time-invariant system on an eight-dimensional latent space with contracting rotational dynamics, a full-rank anisotropic control map, and a set of learned contents. A fixed non-isometric readout assigns each latent state a location in a three-dimensional represented "physical" space, so that represented distance can be rescaled independently of the dynamics. A transition is driven by minimum-energy optimal control, and its latency is the first horizon by which the target is reachable under a fixed control-energy budget. The model is illustrative. It instantiates the definitions above, and it is not fit to neural recordings; every number it produces is a key in the committed `results.json`, and the computation is deterministic given its recorded seed.

The rescaling argument holds inside the model without slack. For a representative pair the latency is 0.642 in model time units, and recomputing it while scaling the represented-space readout by factors from 1 to 10^{20} leaves that latency bit-identical, with a maximum relative deviation of 0, while the apparent velocity $D_{\text{represented}}/\tau$ climbs across 20 orders of magnitude with a log-log slope of 1.000 (Figure 1). The readout scale never enters the transition dynamics, so it cannot change a physical rate. This is the impossibility argument of the earlier section, made visible.

The relevant distance is the control-geometric one. Across 397 random content pairs, the latency is well predicted by the control energy the transition needs at the ensemble's characteristic timescale, with $R^2 = 0.87$ between their logarithms, and it is essentially unrelated to the represented physical distance, with $R^2 = 0.02$ (Figure 2). A naive Euclidean distance in the latent coordinates does little better than the represented distance, at $R^2 = 0.04$, which is the point sharpened: the metric is not incidental, and the control-effort geometry is what tracks how long a transition takes, where a raw distance, whether neural or represented, does not. The represented distance a naive observer would reach for says almost nothing about the timing.

Jump and traverse dissociate as predicted, and forced speed costs fidelity by an amount the neural path sets. Holding the endpoints of a transition fixed and requiring the traversal to visit g intermediate waypoints, the jump latency is flat in g, with a slope of 0, while the traversal latency is affine in g with a mean slope of 0.147 and an affine fit of $R^2 = 0.95$, so that at ten waypoints the traversal

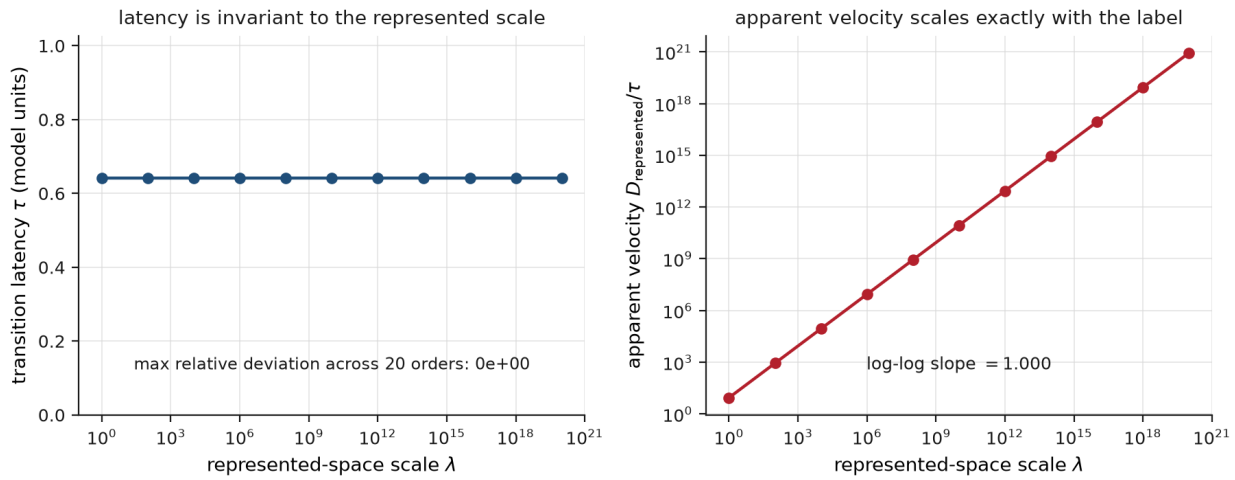

Figure 1: Latency is invariant to the represented scale, while the apparent velocity is not. Left: the transition latency τ for a fixed content pair, recomputed as the represented-space readout is scaled by λ from 1 to 10^{20}; it does not move (maximum relative deviation 0). Right: the apparent velocity $D_{\text{represented}}/\tau$ for the same transition, rising in proportion to λ (log-log slope 1.000). The physical process is one process; the "speed" is a property of how the represented world is labelled.

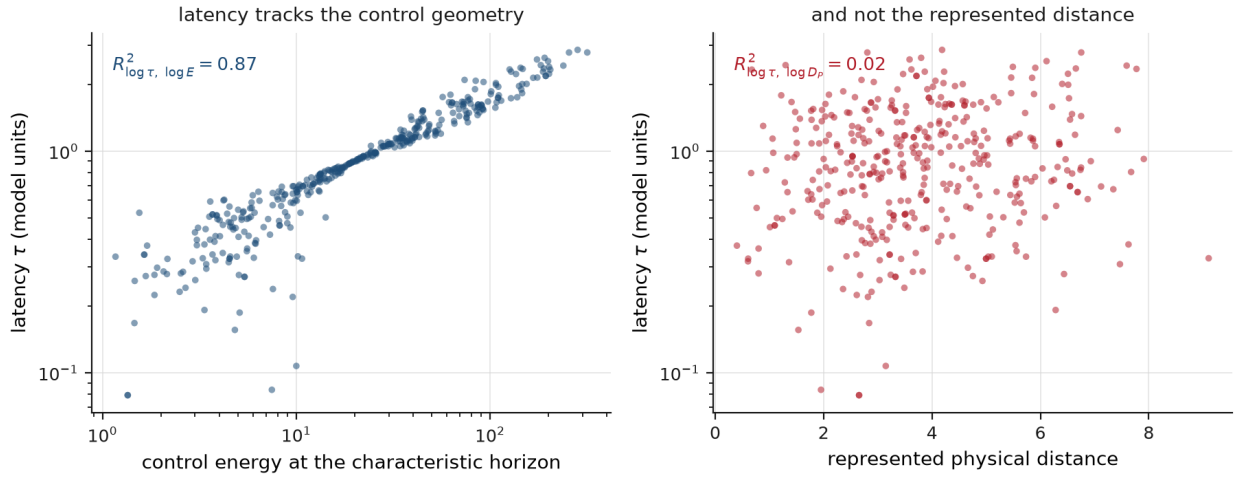

Figure 2: Latency tracks the control geometry rather than the represented distance. Left: log-log relation between transition latency and the control energy required at the characteristic horizon, across 397 content pairs ($R^2 = 0.87$). Right: the same latencies against represented physical distance under the fixed non-isometric readout ($R^2 = 0.02$), a flat cloud. A naive Euclidean latent distance fares no better than the represented one ($R^2 = 0.04$); only the control-effort metric predicts timing.

costs 2.4 times the jump for the same endpoints (Figure 3, left). Demanding a transition at 0.6 of its budget latency leaves the trajectory short of the target, and the resulting fidelity correlates with the neural path length at -0.585 and with the represented distance at only -0.211 (Figure 3, right): a fast transition can still be faithful when the neural path is efficient, and represented distance again says little about which transitions pay the price.

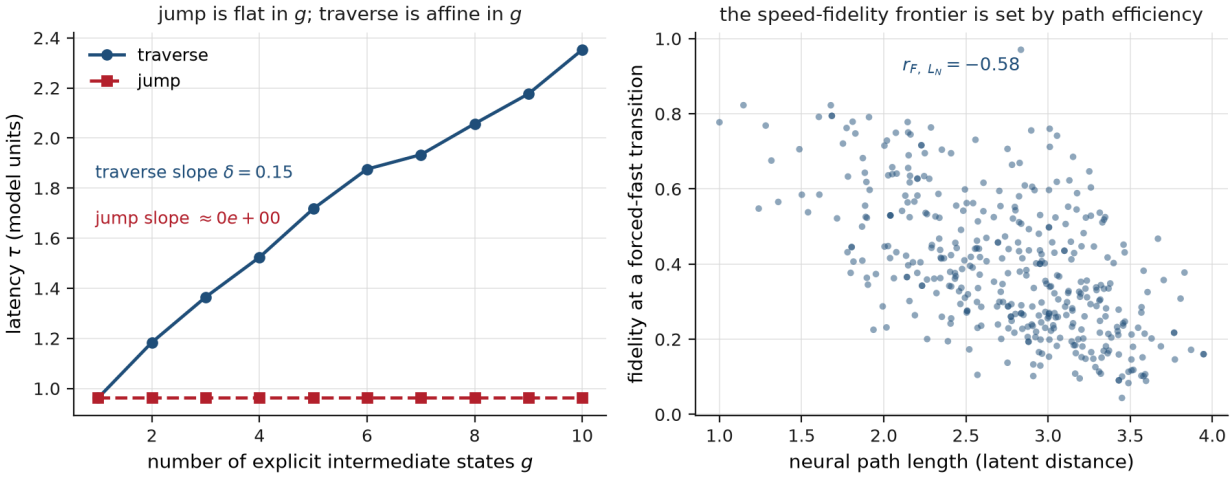

Figure 3: Jump and traverse dissociate, and the speed-fidelity frontier follows the neural path. Left: for fixed endpoints, jump latency is flat in the number of explicit intermediate states g (slope 0), while traverse latency is affine in g (slope $\delta = 0.147$, affine $R^2 = 0.95$), reaching 2.4 times the jump cost at ten waypoints. Right: fidelity when each transition is forced to 0.6 of its budget latency, against neural path length; efficient paths keep fidelity under speed pressure ($r = -0.585$), and represented distance predicts the loss far less ($r = -0.211$).

None of these facts is a claim about the brain. They are properties of a control model built to the framework's definitions, and their role is to show that the definitions are consistent and to fix what the corresponding measurements would look like. The measurements themselves are the experiment.

An experiment that separates the two distances

The design turns on making represented distance and neural distance vary independently, which the world does not do for free. Construct an artificial universe of 16 to 64 visually distinct scenes and, for each pair, manipulate the represented physical distance D_P, the visual similarity D_V, the semantic similarity D_S, and the learned graph distance D_G orthogonally, with familiarity and visual complexity matched. Assign the same relational environment different physical scales across groups, a metre between neighbours for one and a million light-years for another, which is the rescaling argument in participants: if the neural trajectories and latencies match, the apparent external velocity differs by eighteen orders of magnitude while nothing in the imaginative dynamics does.

Within that universe, each trial cues a target and instructs one of three policies, jump directly to

the target, traverse the learned path to it, or continuously morph the current scene into it, followed by an objective probe of a detail of the target so that a verbal label cannot stand in for a constructed image. Decoding uses perception-trained classifiers, which detect imagined content earlier than imagery-trained ones but with an onset that depends on signal-to-noise ratio and trial variability (Dijkstra et al., 2018), so onset is estimated with a change-point model rather than a single threshold, and reported as an estimated representational onset rather than the instant experience begins. Traveling waves are estimated by two independent methods, Hilbert phase gradients and optical flow, with source reconstruction, spatial-autocorrelation-preserving surrogates, and eye-movement controls, because scalp measurements can manufacture wave speed (Orczyk & Kajikawa, 2021); intracranial recordings would settle the wave measurements most convincingly.

The statistics are a hierarchical model of log latency,

$$\log \tau_{ij} = \beta_0 + \beta_P \log D_{P,ij} + \beta_S D_{S,ij} + \beta_V D_{V,ij} + \beta_G D_{G,ij} + \beta_C C_{ij} + \beta_F F_{ij} + \gamma^\top W_{ij} + u_{\text{part}} + u_{\text{item}} + \epsilon_{ij},$$

compared as a nested sequence: M_0 with behavioural variables, M_1 adding neural state-space variables, M_2 adding traveling-wave variables. The wave account is justified only if M_2 improves out-of-sample prediction over M_1; a wave term that adds nothing beyond population activity and cue identity would weaken the strong form of the hypothesis, which is the outcome the design is built to be able to report. The predictions are specific. A direct jump has $\beta_P \approx 0$. Graph distance rather than represented physical distance predicts traversal duration. Neural path length predicts latency better than represented physical distance. Wave topology, coherence, or direction predicts the residual latency and the fidelity. And extremely short transitions sometimes lose fidelity unless the neural path is efficient, the frontier the model traces.

One prediction is sharp enough to stand alone. If externally driven vision runs feedforward waves and internally generated signals run feedback waves, then perception and imagery should propagate in opposite directions,

$$\text{perception} \Rightarrow \text{posterior-to-anterior}, \qquad \text{imagery} \Rightarrow \text{anterior-to-posterior},$$

with the imagery cascade running from hippocampal and prefrontal retrieval through backward alpha or beta waves that reorganize visual excitability to target-specific activity in sensory cortex. The directional asymmetry has been shown for visual stimulation against rest, where stimulation drives forward waves and rest drives backward ones (Alamia & VanRullen, 2019), consistent with feedback carrying predictions in the predictive-coding sense (Rao & Ballard, 1999). That study did not test deliberate imagery, which is what makes imagery a clean new case: the framework predicts that faster and more accurate imagery trials show earlier and more coherent backward organization, wave arrival in sensory cortex preceding target decoding, and state-space trajectories that reach the target basin more directly, while less efficient trials show longer paths, reversals, and unstable wave topologies. A failure of the reversal to appear would put pressure on the predictive-

coding reading rather than on the kinematic framework, and the two claims are worth keeping separate.

Why a representational jump is not a superluminal signal

Special relativity constrains the propagation of signals, energy, and causal influence, with c the invariant it is built around (Einstein, 1905). A genuinely faster-than-light mental process would have to do one of three things: obtain presently unknown information from a remote location faster than light could bring it, transmit a controllable message to that location, or influence an event outside the brain's future light cone. Imagining the sealed card on Mars does none of them. The represented distance is not an interval that anything traverses, the represented magnitude is a variable the local dynamics reset, and the substrate that does the resetting runs at a few metres per second (Van Blooijs et al., 2023), eight orders of magnitude below light. Reference to a place is not a channel to it, and the apparent velocity is the artifact the first section named.

What survives is more useful than the paradox that motivated it. The distance that governs an imaginative transition is the neural one, measured under the control-effort metric, together with the control energy and the number of intermediate states a policy chooses to instantiate, and the simulation shows these predicting latency and fidelity while the represented physical distance predicts neither. The speed of imagination resolves into a profile of separately measurable rates and a mechanism, cortical waves gating the subspaces that content occupies, whose wave-to-state half already has support and whose wave-to-content half the proposed experiment is built to test. Imagination reaches a distant galaxy quickly because local neural dynamics can change what they represent without moving through the represented interval, and how quickly they change it is set by the transition they must make rather than the distance they appear to cross.

References

Alamia, A., & VanRullen, R. (2019). Alpha oscillations and traveling waves: signatures of predictive coding? *PLOS Biology*, 17(10), e3000487.

Batabyal, T., Brincat, S. L., Donoghue, J. A., Lundqvist, M., Mahnke, M. K., & Miller, E. K. (2026). State-space trajectories and traveling waves following distraction. *Journal of Cognitive Neuroscience*, 38(4), 695–715.

Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron*, 100(2), 490–509.

Bhattacharya, S., Brincat, S. L., Lundqvist, M., & Miller, E. K. (2022). Traveling waves in the prefrontal cortex during working memory. *PLOS Computational Biology*, 18(1), e1009827.

Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., & Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature*, 487(7405), 51–56.

Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory* (2nd ed.). Hoboken, NJ: Wiley-Interscience.

Das, A., Zabeh, E., Ermentrout, B., & Jacobs, J. (2026). Planar, spiral, and concentric traveling waves distinguish behavioral states in human memory. *Nature Communications*, 17, 5143.

Dijkstra, N., Mostert, P., de Lange, F. P., Bosch, S., & van Gerven, M. A. J. (2018). Differential temporal dynamics during visual imagery and perception. *eLife*, 7, e33904.

Einstein, A. (1905). Zur Elektrodynamik bewegter Körper. *Annalen der Physik*, 322(10), 891–921.

Gallego, J. A., Perich, M. G., Miller, L. E., & Solla, S. A. (2017). Neural manifolds for the control of movement. *Neuron*, 94(5), 978–984.

Gu, S., Pasqualetti, F., Cieslak, M., Telesford, Q. K., Yu, A. B., Kahn, A. E., Medaglia, J. D., Vettel, J. M., Miller, M. B., Grafton, S. T., & Bassett, D. S. (2015). Controllability of structural brain networks. *Nature Communications*, 6, 8414.

Huang, Q., Xiao, Z., Yu, Q., Luo, Y., Xu, J., Qu, Y., Dolan, R. J., Behrens, T. E. J., & Liu, Y. (2024). Replay-triggered brain-wide activation in humans. *Nature Communications*, 15, 7185.

Kosslyn, S. M., Ganis, G., & Thompson, W. L. (2001). Neural foundations of imagery. *Nature Reviews Neuroscience*, 2(9), 635–642.

Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis: connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2, 4.

Liu, Y., Dolan, R. J., Kurth-Nelson, Z., & Behrens, T. E. J. (2019). Human replay spontaneously reorganizes experience. *Cell*, 178(3), 640–652.

Miller, E. K., Brincat, S. L., & Roy, J. E. (2024). Cognition is an emergent property. *Current Opinion in Behavioral Sciences*, 57, 101388.

Muller, L., Chavane, F., Reynolds, J., & Sejnowski, T. J. (2018). Cortical travelling waves: mechanisms and computational principles. *Nature Reviews Neuroscience*, 19(5), 255–268.

Orczyk, J. J., & Kajikawa, Y. (2021). Magnifying traveling waves on the scalp. *Brain Topography*, 35(1), 162–168.

Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87.

Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171(3972), 701–703.

Van Blooijs, D., van den Boom, M. A., van der Aar, J. F., Huiskamp, G. M., Castegnaro, G., Demuru, M., Zweiphenning, W. J. E. M., van Eijsden, P., Miller, K. J., Leijten, F. S. S., & Hermes, D. (2023). Developmental trajectory of transmission speed in the human brain. *Nature Neuroscience*, 26(4), 537–541.