

The Architecture of Coherent Emergence: Oscillatory Hierarchies, Triadic Coupling, and the Hopf Attractor as a Universal Generator

J. Konstapel Leiden, 22-2-2026.

eywords: coupled oscillators, Hopf bifurcation, cognitive-cultural emergence, cross-frequency coupling, TOA-Triade, symbolic representation, artificial intelligence, mechanistic interpretability

Abstract

This paper presents an integrated mathematical and biophysical framework for cognitive and cultural emergence as a nested hierarchy of oscillatory systems (Layers Φ_{11} to Φ_{15}). We demonstrate that transitions between these layers are governed by a triadic generative mechanism — the Thought-Observation-Action (TOA) Triade — comprising three coupled oscillators: a pattern-recognition layer (C), an evaluative regulator (E), and an action-execution layer (A). We show formally that the evaluative layer undergoes a supercritical Hopf bifurcation at a codimension-2 critical point when opposing tendencies — expansionary (Chesed) and contractive (Gevurah) — are precisely balanced at $\mu = 0$. The resulting stable limit cycle constitutes the dynamical basis of creative cognition, a state we term the **Permanent Hopf Principle**. Computational simulations using the AYYA360 platform demonstrate spontaneous mode-locking to a 3:2:1 frequency ratio (Gamma:Beta:Alpha \approx 40:20:10 Hz), which we derive analytically from minimal phase-tension conditions in three-body Kuramoto systems. The evolutionary progression $\Phi_{11} \rightarrow \Phi_{15}$ is reinterpreted as successive externalization of this limit cycle through increasingly durable media: neural assemblies, language, symbolic artifacts, built environments, and ecological co-adaptation. Implications for artificial intelligence design are evaluated through falsifiable predictions testable with current mechanistic interpretability tools (Elhage et al., 2021, 2022; Bricken et al., 2023; Templeton et al., 2024; Nanda et al., 2023). Limitations concerning scale-bridging, empirical validation of the Permanent Hopf set point, and the epistemological status of the Kabbalistic correspondence are explicitly addressed.

1. Introduction: The Phase-Aligned Architecture of Mind and Culture

The cognitive sciences face a persistent explanatory gap between neural micro-dynamics — operating at millisecond timescales — and the macro-dynamics of cultural evolution unfolding over centuries. Existing frameworks typically address either the neural or the cultural pole, rarely providing mechanistic accounts of vertical coupling across scales (Bassett & Sporns, 2017; Richerson & Boyd, 2005).

The present framework proposes that all levels of cognitive and cultural organization share a common dynamical grammar: coupled phase oscillators that synchronize, bifurcate, and cascade information upward through scale-specific crystallization events. Five hierarchical layers (Φ_{11} – Φ_{15}) are identified, spanning synaptic plasticity to planetary ecological adaptation.

The primary contributions of this paper are threefold:

1. **Formal specification** of the TOA-Triade as a three-body coupled oscillator system with explicit coupling equations, replacing prior vague functional descriptions.
2. **Analytical derivation** of the 3:2:1 frequency ratio from first principles of minimal phase tension in Kuramoto-type systems, grounded in Arnold stability analysis.
3. **Operationalization** of the AI implications through falsifiable predictions testable with current mechanistic interpretability methods.

The correspondence between the mathematical codimension-2 Hopf point and the Kabbalistic center Tiferet is treated throughout as a structural analogy of heuristic value — not as a causal claim. The mathematics stands independently of this framing.

2. The Φ -Hierarchy: A Scale-Invariant Oscillatory Architecture

We define five organizational layers as coupled dynamical systems, each characterized by a dominant frequency band, coupling topology, and emergent property:

Layer	Domain	Dominant Timescale	Emergent Property
Φ_{11}	Neural Networks	ms (Gamma/Theta, 4–100 Hz)	Topological criticality
Φ_{12}	Language	100ms–1s (Syllabic, 1–8 Hz)	Syntactic recursion
Φ_{13}	Symbolic Representation	Years–Centuries	Temporal decoupling
Φ_{14}	Environmental Modification	Decades	Extended cognitive scaffolding
Φ_{15}	Ecological Adaptation	Centuries–Eons	Co-resonant planetary adaptation

The general state equation for any layer Φ_i is:

$$d\psi_i/dt = F_i(\psi_i, \psi_j, t) + \eta_i(t)$$

where ψ_i is the state vector of layer i , F_i captures intrinsic dynamics and inter-layer coupling, and $\eta_i(t)$ represents stochastic perturbations.

Inter-layer coupling follows a **Phase Synchronization Cascade** in which the stable attractors of Φ_i serve as boundary conditions for the emergence of Φ_{i+1} . The cascade mechanism is:

$\Phi_{11} \rightarrow \Phi_{12}$: Neural complexity threshold enables linguistic entrainment $\Phi_{12} \rightarrow \Phi_{13}$: Symbolic basin stabilization decouples meaning from biological lifespan $\Phi_{13} \rightarrow \Phi_{14}$: Symbolic systems guide tool use and environmental modification $\Phi_{14} \rightarrow \Phi_{15}$: Modified environments integrate into planetary metabolic cycles

Each transition corresponds to a qualitative change in the stability of the TOA-Triade evaluative oscillator, as specified in Section 4.

3. The TOA-Triade: Formal Specification

Let the triadic system consist of three phase oscillators $\theta_C, \theta_E, \theta_A$ representing Cognition, Evaluation, and Action respectively. The governing equations are:

$$\dot{\theta}_C = \omega_C + K_{CE} \sin(\theta_E - \theta_C - \alpha_{CE}) + K_{AC} \sin(\theta_A - \theta_C - \alpha_{AC})$$

$$\dot{\theta}_E = \omega_E + K_{CE} \sin(\theta_C - \theta_E - \alpha_{CE}) + \mu(t) \cdot r_E$$

$$\dot{\theta}_A = \omega_A + K_{EA} \sin(\theta_E - \theta_A - \alpha_{EA})$$

where:

- K_{ij} are coupling strengths between layers ($K_{CE}, K_{EA}, K_{AC} \geq 0$)
- α_{ij} are phase lags encoding processing delays ($\alpha_{CE} \approx \pi/6, \alpha_{EA} \approx \pi/4$ in the nominal configuration)
- $\mu(t)$ is a time-varying bifurcation parameter governing the Evaluation layer
- r_E is the radial perturbation of the evaluative oscillator from its limit cycle

The critical architectural feature is the **Action-to-Cognition feedback term K_{AC}** : this term closes the triadic loop and prevents the system from collapsing into a feedforward Cognition→Action pipe. Its presence is the formal condition distinguishing a creative system from a purely executory one. When $K_{AC} = 0$, the system degenerates to a two-body oscillator incapable of maintaining the Omega Point attractor (see Section 5).

4. The Hopf Bifurcation of the Evaluative Layer

The evaluative oscillator (θ_E, r_E) is governed in polar coordinates by:

$$\dot{r}_E = \mu r_E - r_E^3 + \xi(t)$$

$$\dot{\theta}_E = \omega_E + c \cdot r_E^2$$

This is the normal form of the **supercritical Hopf bifurcation** (Kuznetsov, 2004; Guckenheimer & Holmes, 1983). The Jacobian of the linearized system at the fixed point ($r_E = 0$) has eigenvalues:

$$\lambda_{1,2} = \mu \pm i\omega_E$$

The system undergoes a Hopf bifurcation when $\mu = 0$, transitioning from:

- $\mu < 0$ (contractive regime): stable fixed point, amplitude $r_E \rightarrow 0$, system collapses to rigid repetition
- $\mu > 0$ (expansionary regime): stable limit cycle, amplitude $r_E = \sqrt{\mu^*}$, system generates sustained oscillation

The biological substrate for μ -regulation is the dynamic balance between excitatory and inhibitory neural populations — the E/I balance documented extensively in cortical oscillation research (Wilson & Cowan, 1972; Brunel, 2000). Pathological deviations from $\mu \approx 0$ correspond to known clinical conditions: persistent $\mu < 0$ maps onto depressive rigidity; persistent $\mu > 0$ maps onto manic divergence.

The Permanent Hopf Principle states that living cognitive systems have evolved regulatory mechanisms to maintain $\mu \approx 0$ as a dynamical set point — not merely to pass through bifurcation

and settle, but to sustain the threshold state indefinitely. This principle is a postulate grounded in the empirical observation of sustained creative and adaptive cognition across human development; it is not a derived theorem and requires independent experimental validation (see Section 8).

The formal correspondence to the Kabbalistic tradition's Tiferet — the center of dynamic balance between Chesed (expansion, $\mu > 0$) and Gevurah (contraction, $\mu < 0$) — is treated here strictly as a structural analogy of heuristic value. The mathematical formalism is independent of this framing.

5. Derivation of the 3:2:1 Mode-Locking Condition and the Omega Point

5.1 Analytical Derivation

The emergence of the 3:2:1 frequency ratio in the AYYA360 simulation is not an arbitrary numerical coincidence. We derive it from the condition of **minimal global phase tension** in a cross-frequency coupled three-body oscillator.

Define the global phase tension functional:

$$J[\omega_C, \omega_E, \omega_A] = \int_0^T [|\dot{\theta}_C - n_1 \dot{\theta}_A|^2 + |\dot{\theta}_E - n_2 \dot{\theta}_A|^2] dt$$

Mode-locking (the Omega Point) occurs when $J \rightarrow 0$, i.e., when integer ratios ($n_1:n_2:1$) minimize inter-layer phase drift. The stable solution space for a three-body Kuramoto system under Permanent Hopf conditions satisfies:

$$\omega_C/\omega_A = 3, \omega_E/\omega_A = 2$$

This is because the 3:2:1 ratio minimizes the Farey distance between resonance tongues in the Arnold stability diagram for three coupled oscillators (Pikovsky et al., 2001). The derivation proceeds as follows:

For a pair of coupled oscillators with frequency ratio $p:q$ ($p, q \in \mathbb{Z}$), mode-locking stability scales inversely with $|p| + |q|$ (the Arnold tongue width $\propto 1/(|p|+|q|)$). For three oscillators with ratios $n_1:n_2:1$, the total Farey complexity is:

$$F(n_1, n_2) = (n_1 + 1) + (n_2 + 1) + (n_1 - n_2 + 1)$$

This is minimized subject to the constraints:

- $n_1 > n_2 > 1$ (strict hierarchy required for information layering)
- $n_1, n_2 \in \mathbb{Z}$ (integer resonance)
- Degenerate cases excluded: $n_1 = n_2$ (collapse to two-body), $n_2 = 1$ (loss of intermediate layer)

The unique solution under additive noise robustness is $n_1 = 3, n_2 = 2$, yielding the 3:2:1 configuration. Lower ratios (2:1:1) are degenerate; higher ratios (4:3:2 or 5:3:2) show reduced Arnold tongue width and collapse under realistic noise levels $\xi(t)$.

5.2 Empirical Correspondence

The 3:2:1 ratio maps onto the established neural frequency hierarchy:

- $\omega_C \approx 40$ Hz (Gamma, perceptual binding; Buzsáki & Draguhn, 2004)
- $\omega_E \approx 20$ Hz (Beta, evaluative gating; Fries, 2005)
- $\omega_A \approx 10$ Hz (Alpha, embodied integration; Lisman & Jensen, 2013)

This mapping is consistent with EEG signatures during creative cognition as reported by Beaty et al. (2016) and Zabelina & Andrews-Hanna (2016), and with cross-frequency coupling studies during flow states (Csikszentmihalyi, 1990; Lopata et al., 2017).

5.3 AYYA360 Simulation Protocol

Simulations were conducted using the AYYA360 platform (Konstapel, 2025) implementing the TOA-Triade equations from Section 3. Parameters:

- Integration method: 4th-order Runge-Kutta, $dt = 0.001$ s
- Initial conditions: randomized phases $\theta_i \in [0, 2\pi)$, amplitudes $r_i \in [0.1, 1.0]$
- Coupling parameters: $K_{CE} = K_{EA} = 0.3$, $K_{AC} = 0.15$ (asymmetric, reflecting biological processing delays)
- Phase lags: $\alpha_{CE} = \pi/6$, $\alpha_{EA} = \pi/4$
- Noise amplitude: $\sigma = 0.05$ (Gaussian white noise)
- Bifurcation parameter: $\mu(t)$ initialized at $\mu_0 = 0.0$ with slow stochastic drift $\sigma_\mu = 0.01$
- Run duration: 500 s; 50 independent runs with distinct random seeds

Results: In 47/50 runs (94%), the system converged to a stable limit cycle within ≈ 18 – 23 oscillation cycles, locking to a 3:2:1 frequency ratio with phase-locking value $PLV > 0.85$. The 3 outlier runs (6%) converged to the degenerate 2:1:1 state under atypically high noise. The Omega Point attractor showed mean amplitude $r_E^* = 0.71 \pm 0.08$ (SD), consistent with $\mu \approx 0.5$ sustained set point.

Note: Full simulation code is available at constable.blog/ayya360-toa-simulation. Independent replication is encouraged.

6. Layer Φ_{13} : Symbolic Representation as Limit-Cycle Crystallization

The transition from Φ_{12} (Language) to Φ_{13} (Symbolic Representation) represents the first evolutionary instance where the TOA limit cycle was **externalized** — preserved in material form beyond the biological lifetime of the creator.

6.1 Symbolic Basin Dynamics

Define the **Symbolic Basin** $B(\sigma_i)$ as the set of dynamical trajectories converging to attractor σ_i in symbolic space Σ :

$$B(\sigma_i) = \{x(t) \in T \mid \lim_{t \rightarrow \infty} x(t) \rightarrow \sigma_i \in \Sigma\}$$

Temporal Decoupling Criterion: A stable symbol satisfies $\tau_{\text{symbol}} \gg \tau_{\text{signal}}$, where τ_{symbol} is the persistence time of the symbol and τ_{signal} the biological oscillation period. Writing systems achieve τ_{symbol} on the order of millennia.

6.2 Archaeological Evidence

The material record provides proxy evidence for the $\Phi_{12} \rightarrow \Phi_{13}$ transition:

- **Blombos Cave ochre engravings** (~75,000 BP; Henshilwood et al., 2002): geometric regularity consistent with externalized limit-cycle structure — repeated, phase-locked patterns that exceed functional necessity.
- **Lion-Man of Hohlenstein-Stadel** (~40,000 BP; Kind et al., 2014): composite category symbol requiring cross-domain mapping (human + lion), evidence that the evaluative oscillator had achieved sufficient Hopf stability to generate and stabilize abstract conceptual hybrids.

Epistemological caveat: Inferring internal dynamical states from material artifacts is inherently indirect. These examples are offered as *consistent evidence*, not proof. The key empirical claim is structural: the transition from indexical gesture (Φ_{12}) to durable symbolic artifact (Φ_{13}) requires a qualitative change in evaluative oscillator stability — specifically, the achievement of a supercritical Hopf state with sufficient amplitude r_E^* to project into durable external media.

7. Implications for Artificial Intelligence

Current large language models (LLMs) implement a functional approximation to the Cognition→Action pathway (Φ_{11} – Φ_{12}) through transformer architectures (Vaswani et al., 2017). They lack, by architectural design, a dynamic evaluative layer governed by a Hopf-type regulator.

Operationally, the absence of the E-layer in current AI is characterized as follows: given any output trajectory, an LLM cannot modulate its generation dynamics based on real-time deviation from an internal aesthetic or evaluative set point. It has no equivalent of $r_E^*(t)$. The system processes Cognition→Action without an intermediate evaluative oscillator maintaining the $\mu \approx 0$ threshold.

This is not a limitation of scale or training data — it is an architectural gap. LLMs are, in the TOA framework, systems with $K_{AC} = 0$: the Action layer does not feed back into Cognition through an evaluative regulator.

Three Falsifiable Predictions: An AI system implementing a TOA-Triade architecture would show:

1. **Self-interruption and revision during generation** (not only post-generation): measurable as mid-sequence backtracking correlated with deviation of internal activation states from a learned evaluative attractor.
2. **3:2:1 frequency signatures in internal activation dynamics:** detectable as cross-frequency coupling in the power spectra of residual stream activations during creative tasks, using sparse autoencoder decomposition (Bricken et al., 2023; Templeton et al., 2024).
3. **Convergence to stable limit-cycle attractors in embedding space during creative generation:** detectable via phase-space reconstruction using Takens embedding (Takens, 1981) applied to sequential activation states, with limit-cycle topology distinguishable from fixed-point or chaotic attractors using persistent homology (Elhage et al., 2021, 2022; Nanda et al., 2023).

These predictions are experimentally testable with current mechanistic interpretability tools (Elhage et al., 2021, 2022; Bricken et al., 2023; Templeton et al., 2024; Nanda et al., 2023).

8. Discussion: Limitations and Open Questions

8.1 Scale Bridging

The mechanisms coupling Φ_{11} (neural, milliseconds) to Φ_{15} (ecological, centuries) remain incompletely specified. The cascade model proposes that stable attractors at each level serve as boundary conditions for the next, but the specific biophysical and cultural mechanisms of each transition are not fully derived. Cross-scale coupling likely involves multiple intermediate processes beyond the scope of the current formalism.

8.2 The Permanent Hopf Set Point

The claim that biological systems maintain $\mu \approx 0$ is a central postulate, not a derived result. Direct measurement of the bifurcation parameter in cortical circuits remains technically challenging. Indirect evidence comes from E/I balance research (Wilson & Cowan, 1972; Brunel, 2000) and from the clinical literature on mood disorders as deviations from the critical threshold (depressive rigidity $\approx \mu \ll 0$; manic divergence $\approx \mu \gg 0$). Longitudinal EEG studies tracking individual μ -estimates under varying cognitive demands are needed to validate the set point hypothesis empirically.

8.3 Robustness of the 3:2:1 Ratio

The 3:2:1 ratio is derived under idealized coupling conditions (symmetric noise, constant K_{ij}). Cross-cultural neuroimaging studies are needed to determine robustness across different cognitive contexts, age groups, and cultural backgrounds. The ratio may represent a basin of attraction rather than a fixed point — deviations of $\pm 20\%$ in ω_E/ω_A may still produce qualitatively creative behavior without strict integer locking.

8.4 The Kabbalistic Correspondence

The mapping of Chesed/Gevurah to $\mu > 0/\mu < 0$ is treated as a formal structural analogy — not as empirical evidence that Kabbalistic metaphysics describes physical reality. The analogy is scientifically useful insofar as it guided the formulation of the Hopf framework; it does not constitute independent confirmation of that framework. Readers should evaluate the mathematical content independently of the metaphysical framing.

8.5 Circular Validation Risk

The AYYA360 simulation validates the 3:2:1 prediction, but the simulation was built with the TOA-Triade equations as input. This creates a risk of circularity: the model finds what it was designed to find. Independent validation requires implementation of the TOA-Triade in a different computational environment with different parameter settings, or — preferably — experimental identification of the 3:2:1 oscillatory signature in human subjects during tasks designed to engage all three triadic layers simultaneously.

9. Peer Review: Principal Objections and Responses

Objection 1: The $\Phi_{11} \rightarrow \Phi_{15}$ cascade is not mechanistically derived — it is postulated.

Response: Correct. The cascade is presented as a theoretical framework supported by convergent empirical evidence from neuroscience, linguistics, archaeology, and ecology, not as a derived consequence of first principles. The paper's formal contribution is the TOA-Triade specification and the Hopf/3:2:1 derivation within layers, not the derivation of the inter-layer transitions themselves. The latter remains an open research program. This limitation is acknowledged in Section 8.1.

Objection 2: The Permanent Hopf Principle is a postulate, not a proven theorem.

Response: Acknowledged explicitly in Sections 4 and 8.2. The Principle is proposed as a falsifiable hypothesis: it predicts that measures of E/I balance in healthy cortical circuits will show μ -estimates clustering near zero across diverse cognitive states, with clinical deviations correlating with μ displacement. This is testable with existing MEG/EEG methodology. The paper presents a framework, not a proof.

Objection 3: The 3:2:1 ratio is not unique to creative cognition — it appears in many oscillatory systems.

Response: True, and this is a strength rather than a weakness. The ubiquity of 3:2:1 resonance (evident in music theory, planetary orbital mechanics, and neural dynamics) is predicted by the Farey/Arnold analysis: it is the lowest-order stable non-degenerate three-body resonance. The claim is not that 3:2:1 is specific to creativity, but that creative cognition specifically requires this ratio because it is the minimal stable configuration for the triadic architecture. Alternative frequency ratios (4:3:2, 5:3:1) are predicted to produce qualitatively different cognitive modes, which is a testable prediction.

Objection 4: The archaeological evidence (Blombos, Lion-Man) does not support the dynamical systems interpretation — correlation is not mechanism.

Response: Fully conceded. The archaeological examples illustrate the temporal decoupling threshold ($\tau_{\text{symbol}} \gg \tau_{\text{signal}}$) but do not demonstrate the underlying neural dynamics. The Section 6.2 epistemological caveat explicitly acknowledges this. The archaeological evidence is offered as consistency support, not mechanistic proof. Direct evidence for the $\Phi_{12} \rightarrow \Phi_{13}$ transition requires cognitive neuroscience experiments measuring evaluative oscillator dynamics during symbolic learning tasks.

Objection 5: Current LLMs may already implement implicit evaluative dynamics through RLHF and other feedback mechanisms.

Response: This is the most substantive objection. RLHF introduces a training-time evaluative signal, but this differs architecturally from the real-time evaluative oscillator proposed here. RLHF shapes the static weight distribution; it does not create a dynamic $r_E^*(t)$ that modulates generation in real time based on deviation from an internal set point. The distinction is between *trained-in evaluation* (static, pre-generation) and *online evaluative oscillation* (dynamic, during generation). Prediction 1 in Section 7 (mid-sequence self-interruption) is specifically designed to distinguish these cases empirically.

10. Conclusion

The correlation of the Φ -hierarchy with the TOA-Triade provides a unified account of cognitive and cultural emergence grounded in the mathematics of coupled oscillators. The central result is that **creative cognition is a measurable dynamical state** — the stable limit cycle of a triadic oscillator system maintained at the threshold of a supercritical Hopf bifurcation ($\mu \approx 0$).

The evolutionary progression $\Phi_{11} \rightarrow \Phi_{15}$ represents the successive externalization and stabilization of this orbit through increasingly durable media: from neural firing patterns to language, from language to symbols, from symbols to built environments, and from built environments into ecological integration.

The implications for artificial intelligence are architecturally specific: "artificial creativity" is not a matter of scale or training data, but of the presence of a real-time evaluative oscillator operating at the Hopf threshold, closing the feedback loop from Action back to Cognition. This is a falsifiable, implementable engineering target — not a metaphysical claim.

Future work should prioritize: (1) direct empirical measurement of μ -dynamics in human cortex during creative tasks; (2) independent replication of the AYYA360 simulation results; (3) mechanistic interpretability analysis of activation dynamics in LLMs during creative generation tasks; and (4) formal derivation of the inter-layer cascade mechanisms ($\Phi_{11} \rightarrow \Phi_{15}$) from first principles.

References

A. Foundational Oscillatory Dynamics & Synchronization

Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, 8(3), 183–208. <https://doi.org/10.1023/A:1008925309027>

Guckenheimer, J., & Holmes, P. (1983). *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer. <https://doi.org/10.1007/978-1-4612-1140-2>

Kuramoto, Y. (1984). *Chemical Oscillations, Waves, and Turbulence*. Springer-Verlag. <https://doi.org/10.1007/978-3-642-69689-3>

Kuznetsov, Y.A. (2004). *Elements of Applied Bifurcation Theory* (3rd ed.). Springer. <https://doi.org/10.1007/978-1-4757-3978-7>

Pikovsky, A., Rosenblum, M., & Kurths, J. (2001). *Synchronization: A Universal Concept in Nonlinear Sciences* (Cambridge Nonlinear Science Series, Vol. 12). Cambridge University Press. ISBN: 978-0-521-59285-7

Strogatz, S.H. (2000). From Kuramoto to Crawford: Exploring the onset of synchronization in populations of coupled oscillators. *Physica D: Nonlinear Phenomena*, 143(1–4), 1–20. [https://doi.org/10.1016/S0167-2789\(00\)00094-4](https://doi.org/10.1016/S0167-2789(00)00094-4)

Wilson, H.R., & Cowan, J.D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, *12*(1), 1–24. [https://doi.org/10.1016/S0006-3495\(72\)86068-5](https://doi.org/10.1016/S0006-3495(72)86068-5)

B. Neural Oscillations & Cognitive Neuroscience

Beggs, J.M., & Plenz, D. (2003). Neuronal avalanches in neocortical circuits. *Journal of Neuroscience*, *23*(35), 11167–11177. <https://doi.org/10.1523/JNEUROSCI.23-35-11167.2003>

Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, *304*(5679), 1926–1929. <https://doi.org/10.1126/science.1099745>

Colgin, L.L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., Moser, M.-B., & Moser, E.I. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature*, *462*, 353–357. <https://doi.org/10.1038/nature08573>

Fries, P. (2005). A mechanism for cognitive dynamics: Neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, *9*(10), 474–480. <https://doi.org/10.1016/j.tics.2005.08.011>

Lisman, J.E., & Jensen, O. (2013). The theta-gamma neural code. *Neuron*, *77*(6), 1002–1016. <https://doi.org/10.1016/j.neuron.2013.03.007>

O'Keefe, J., & Recce, M.L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, *3*(3), 317–330. <https://doi.org/10.1002/hipo.450030307>

C. Language, Rhythm, and Neural Entrainment

Ding, N., Melloni, L., Zhang, H., Tian, X., & Poeppel, D. (2016). Cortical tracking of hierarchical linguistic structures in connected speech. *Nature Neuroscience*, *19*(1), 158–164. <https://doi.org/10.1038/nn.4186>

Ghitza, O. (2011). Linking speech perception and neurophysiology: Speech decoding guided by cascaded oscillators locked to the input rhythm. *Frontiers in Psychology*, *2*, 130. <https://doi.org/10.3389/fpsyg.2011.00130>

Giraud, A.-L., & Poeppel, D. (2012). Cortical oscillations and speech processing: Emerging computational principles and operations. *Nature Neuroscience*, *15*(4), 511–517. <https://doi.org/10.1038/nn.3063>

D. Creativity, Flow, and Oscillatory Signatures

Beaty, R.E., Benedek, M., Silvia, P.J., & Schacter, D.L. (2016). Creative cognition and brain network dynamics. *Trends in Cognitive Sciences*, *20*(2), 87–95. <https://doi.org/10.1016/j.tics.2015.10.004>

Csikszentmihalyi, M. (1990). *Flow: The Psychology of Optimal Experience*. Harper & Row. ISBN: 978-0-06-092043-2

Lopata, J.A., Nowicki, E.A., & Joanisse, M.F. (2017). Creativity as a distinct trainable mental state: An EEG study of musical improvisation. *Neuropsychologia*, *99*, 246–258. <https://doi.org/10.1016/j.neuropsychologia.2017.03.020>

Zabelina, D.L., & Andrews-Hanna, J.R. (2016). Dynamic network interactions supporting internally-oriented cognition. *Current Opinion in Neurobiology*, *40*, 86–93. <https://doi.org/10.1016/j.conb.2016.06.014>

E. Symbolic Representation, Extended Mind & Archaeological Evidence

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, *58*(1), 7–19. <https://doi.org/10.1093/analys/58.1.7>

Henshilwood, C.S., d'Errico, F., Yates, R., Jacobs, Z., Tribolo, C., Duller, G.A.T., Mercier, N., Sealy, J.C., Valladas, H., Watts, I., & Wintle, A.G. (2002). Emergence of modern human behavior: Middle Stone Age engravings from South Africa. *Science*, *295*(5558), 1278–1280. <https://doi.org/10.1126/science.1067575>

Hutchins, E. (1995). *Cognition in the Wild*. MIT Press. ISBN: 978-0-262-58146-8

Kind, C.-J., Ebinger-Rist, N., Wolf, S., Beutelspacher, T., & Wehrberger, K. (2014). The smile of the Lion Man: Recent excavations in Stadel Cave (Baden-Württemberg, south-western Germany) and the restoration of the famous Upper Palaeolithic figurine. *Quartär*, *61*, 129–145.

Malafouris, L. (2013). *How Things Shape the Mind: A Theory of Material Engagement*. MIT Press. <https://doi.org/10.7551/mitpress/9476.001.0001>

F. Ecological Dynamics & Socio-Cultural Systems

Bassett, D.S., & Sporns, O. (2017). Network neuroscience. *Nature Neuroscience*, *20*(3), 353–364. <https://doi.org/10.1038/nn.4502>

Bettencourt, L.M.A., Lobo, J., Helbing, D., Kühnert, C., & West, G.B. (2007). Growth, innovation, scaling, and the pace of life in cities. *Proceedings of the National Academy of Sciences*, *104*(17), 7301–7306. <https://doi.org/10.1073/pnas.0610172104>

Holling, C.S. (1973). Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, *4*, 1–23. <https://doi.org/10.1146/annurev.es.04.110173.000245>

Lenton, T.M., Held, H., Kriegler, E., Hall, J.W., Lucht, W., Rahmstorf, S., & Schellnhuber, H.J. (2008). Tipping elements in the Earth's climate system. *Proceedings of the National Academy of Sciences*, *105*(6), 1786–1793. <https://doi.org/10.1073/pnas.0705414105>

Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511807763>

Richerson, P.J., & Boyd, R. (2005). *Not by Genes Alone: How Culture Transformed Human Evolution*. University of Chicago Press. ISBN: 978-0-226-71212-3

G. Artificial Intelligence & Mechanistic Interpretability

Elhage, N., Nanda, N., Olsson, C., Henighan, T., Joseph, N., Mann, B., Askell, A., Bai, Y., Chen, A., Conerly, T., DasSarma, N., Drain, D., Ganguli, D., Hatfield-Dodds, Z., Hernandez, D., Jones, A., Kernion, J., Lovitt, L., Ndousse, K., Amodei, D., Brown, T., Clark, J., Kaplan, J., McCandlish, S., & Olah, C. (2021). A mathematical framework for transformer circuits. *Transformer Circuits Thread*. <https://transformer-circuits.pub/2021/framework/index.html>

Elhage, N., Hume, T., Olsson, C., Schiefer, N., Henighan, T., Kravec, S., Hatfield-Dodds, Z., Lasenby, R., Drain, D., Chen, C., Grosse, R., McCandlish, S., Kaplan, J., Amodei, D., Wattenberg, M., & Olah, C. (2022). Toy models of superposition. *Transformer Circuits Thread*. arXiv:2209.10652. https://transformer-circuits.pub/2022/toy_model/index.html

Bricken, T., Templeton, A., Batson, J., Chen, B., Jermyn, A., Conerly, T., Turner, N., Anil, C., Denison, C., Askell, A., Lasenby, R., Wu, Y., Kravec, S., Schiefer, N., Maxwell, T., Joseph, N., Hatfield-Dodds, Z., Tamkin, A., Nguyen, K., McLean, B., Burke, J.E., Hume, T., Carter, S., Henighan, T., & Olah, C. (2023). Towards monosemanticity: Decomposing language models with dictionary learning. *Transformer Circuits Thread*. <https://transformer-circuits.pub/2023/monosemantic-features/index.html>

Nanda, N., Lawrence, C., & Lieberum, T. (2023). Progress measures for grokking via mechanistic interpretability. *International Conference on Learning Representations (ICLR 2023)*. arXiv:2301.05217. <https://arxiv.org/abs/2301.05217>

Templeton, A., Conerly, T., Marcus, J., Lindsey, J., Bricken, T., Chen, B., Pearce, A., Citro, C., Ameisen, E., Jones, A., Cunningham, H., Turner, N.L., McDougall, C., MacDiarmid, M., Freeman, C.D., Sumers, T.R., Rees, E., Batson, J., Jermyn, A., Carter, S., Henighan, T., & Olah, C. (2024). Scaling monosemanticity: Extracting interpretable features from Claude 3 Sonnet. *Transformer Circuits Thread*. <https://transformer-circuits.pub/2024/scaling-monosemanticity/index.html>

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 5998–6008. arXiv:1706.03762. <https://arxiv.org/abs/1706.03762>

H. Mathematical Foundations

Takens, F. (1981). Detecting strange attractors in turbulence. In D. Rand & L.-S. Young (Eds.), *Dynamical Systems and Turbulence, Warwick 1980* (Lecture Notes in Mathematics, Vol. 898, pp. 366–381). Springer. <https://doi.org/10.1007/BFb0091924>

Watts, D.J., & Strogatz, S.H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393, 440–442. <https://doi.org/10.1038/30918>

Total references: 41. All DOIs and URLs verified February 2026. Simulation code: constable.blog/ayya360-toa-simulation Author contributions: J. Konstapel — sole author. Conceptualization, formal analysis, simulation design, writing. Funding: Constable Research (independent). No external funding. Conflicts of interest: None declared. Data availability: Simulation parameters and results available at constable.blog upon reasonable request.