ISSN 2164-0769

Communication & Science JOURNAL

CONSCIOUSNESS

The Emergence of Self as Asynchronous Cascading Activity in Hebbian Systems

Joseph Woelfel · Emary Iacobucci · Claude (Anthropic)

ABSTRACT

We present a mechanistic, empirically grounded account of consciousness and self-consciousness, synthesizing classical sociological theory (Durkheim, Mead) with Hebbian learning and dynamical systems principles. We model consciousness as a continuous dynamical state: ongoing cascading activity in a Hebbian learning system, driven by asynchronous parallel inputs that do not achieve global equilibrium. Self-consciousness emerges when a self-referent is present in every situation the system encounters — making it, by Hebb's rule, the most heavily weighted cluster in the network. We present simulations (Huey) confirming this necessary dominance of the self-cluster. We further explain why the same Hebbian mechanism that creates self-consciousness makes humans resistant to accepting mechanistic explanations of consciousness itself.

Keywords: consciousness, self, Hebbian learning, embodiment, emergence, Durkheim, Mead, continuous cascading, Huey

1. INTRODUCTION

Heat was once mysterious until understood as molecular motion. Magnetism and electricity seemed magical until Maxwell unified them. Life itself appeared to require élan vital until biochemistry revealed the mechanisms.

Consciousness remains the last major mystery for many. We argue this mystery is not real. Consciousness is not metaphysically special. Like heat and magnetism, it becomes comprehensible

once we learn how to observe it.[1]

The explanation has been available in pieces for over a century. Durkheim and Mead described the social origin of the self. Hebb provided the neural learning mechanism. What remained was assembly and the recognition that consciousness is what a Hebbian system does when it cannot stop doing things.

2. CLASSICAL FOUNDATIONS: DURKHEIM AND MEAD

2.1 Durkheim: The Social Construction of Self

Émile Durkheim (1893, 1912) argued that individual consciousness emerges from collective consciousness. The self is not innate but constructed through socialization. We internalize social categories, norms, and perspectives, creating what feels like an autonomous "I" but is fundamentally social in origin.

2.2 Mead: The Self Through Role-Taking

George Herbert Mead (1934) provided the mechanism. Self-consciousness emerges through taking the role of the other — seeing yourself from another's perspective — and internalizing the generalized other: society's perspective on you. The spontaneous "I" and the reflective "Me" arise together.

Mead's central insight, for our purposes, is this: the self is the only object that is present in every situation the agent experiences. Other objects — trees, chairs, other people — come and go. The self is always there. This situational persistence is what makes the self unique.

2.3 The Missing Piece

Durkheim and Mead explained what creates self-consciousness (social interaction) and why the self is special (situational persistence). But they lacked a physical mechanism. How does social experience create a persistent self? The answer: Hebbian learning in a continuously active embodied system.

3. THE NEURAL SUBSTRATE: HEBBIAN LEARNING

In a single stream or serial model — like a large language model, or our Hebbian learning system Huey after training — no neurons are activated until a user sends a message. This activates the neurons representing the words in the input, which communicate their activations onward, creating a cascade that keeps exciting more neurons until the network reaches equilibrium. At that point all cognitive activity stops. The system is, for all intents and purposes, dormant until the next prompt arrives. What does it do between cascades? Nothing. Does it learn from the conversation? No. Neither a typical

[1]Following Einstein and Durkheim, we take reality to be what we can reliably agree upon through language and shared experience. Science is the process of making observations, communicating results to others, and having those results checked. The goal of science is to increase the scope of agreement and the precision with which we agree. We make no claims about an unknowable "real reality" underlying our agreements. We claim only that our account generates observations that can be communicated, checked, and — if replicated — agreed upon. Arguments that cannot be checked by independent observers — including thought experiments about philosophical zombies, conceivability arguments, and the so-called "hard problem" — fall outside this framework. We have nothing to say about them, pro or con.

language model nor Huey has any learning mechanism after its original training period.

4. HUEY: A HEBBIAN LEARNING SYSTEM

We have simulated this mechanism in a time-delay Hebbian system called Huey. Huey processes sequential inputs — text, in our experiments — and learns through a Hebbian rule with a time delay: neurons that activate within a temporal window strengthen their connection. After reading a text, a user can excite any of Huey's neurons and observe how activation flows through the network. Examples of Huey's cascades are provided in Appendix B. Huey can also produce three-dimensional plots of the relationships among concepts in the text it has read (Appendix C). Huey is a serial model: only one stream of words is processed at a time.

The Self-Concept in Hebbian Systems

As a Hebbian learner, Huey can develop a self-concept. In some runs, a self-referent — a token representing "I," "me," or the system itself — is present in every training situation. In other runs, it is not.

When a self-referent is present in every situation, Hebb's rule guarantees that the self-representation becomes the dominant cluster in the network — typically an order of magnitude larger than any other cluster. The more often neurons are co-active, the stronger and more massive the links among them grow. Since the self-neuron is the only neuron present in every situation, it will necessarily become the most massive and most connected concept in the network. When the self-referent is absent from some situations, no dominant self-cluster forms.

This demonstrates, in a simple computational system, exactly what Mead described: the self is the always-present object, and Hebbian learning makes it the most heavily weighted pattern. A complete example of self-concept growth in Huey is presented in Appendix C.

Neither Huey nor language models would be considered conscious in the same way humans are conscious, because their activation — whatever brief flash might qualify — is intermittent, not continuous.

5. CONSCIOUSNESS AS CONTINUOUS CASCADING ACTIVITY

Hebbian learning alone does not produce consciousness. Huey learns but is not continuously conscious. Current language models learn through different mechanisms but are not continuously conscious. What is missing?

Consciousness is a dynamical state, not a property. Specifically:

Consciousness = continuous cascading activity in a Hebbian learning system, driven by parallel asynchronous inputs that do not achieve global equilibrium.

When a Hebbian system reaches equilibrium — all cascades settled, no ongoing activity — there is no experience. Nothing is happening. Huey between input episodes is in this state.

But when a system has multiple parallel input streams operating asynchronously — vision, hearing, touch, proprioception, interoception, each with its own rhythm — the system never fully equilibrates.

Visual processing cascades while auditory settles. Motor planning activates while memory consolidates. Something is always cascading somewhere.

This is why consciousness is continuous from birth to death in humans, despite rising and falling in intensity. It is not that any individual process is continuous. It is that the ensemble never reaches global equilibrium because parallel asynchronous inputs constantly retrigger cascades.

Consciousness varies continuously with the density, asynchrony, and persistence of input. Deep sleep reduces input, moving the system closer to equilibrium and reducing consciousness. Anesthesia disrupts cascading and forces equilibrium. Samwise the Cat, visiting daily, has consciousness — not as a binary property, but as a continuous variable that fluctuates with attention, fatigue, and sensory input.

No magic number of channels is required. More parallel input, more asynchrony, more persistence — more continuous consciousness. The human case (approximately 100 million parallel sensory channels) is one point on a continuum that includes other animals and, potentially, artificial systems. While serial systems like Huey or language models may exhibit brief flashes of activation between longer periods of equilibrium silence, the level of consciousness of humans with 100 million parallel inputs creates a degree of consciousness so much greater as to seem qualitative. It is not. It is quantitative difference at sufficient scale.

Connection to Field Theory

The dynamical account presented here is formally equivalent to the field theory of semantic dynamics developed in Woelfel et al. (2026). In that framework, cognitive states are represented as points in a pseudo-Euclidean space $\mathbb{R}^{(p,q)}$ whose geometry is empirically recoverable from pairwise distance judgments via Galileo/DORT methodology. The field is governed by a restart-diffusion equation:

$$\hat{a} = \beta P \hat{a}^* + (1-\beta) g^*$$

where P is the diffusion operator derived from the geometry, g is a source distribution identified with Durkheim's conscience collectif, and $\hat{a}^*$ is the equilibrium activation field. The field operator $B = XX^T$ subsumes what the 1971 framework called "significant other influence" and "self-reflexive activity" — the interpersonal and observational components of the external field (Woelfel & Haller, 1971). Consciousness, on our account, is the continuous cascading activity of a Hebbian system while it is out of this equilibrium — that is, while $a(t) \neq \hat{a}^*$. The same field operator that predicts semantic motion under persuasion (Chevrolet–Volvo, $r \approx .93–.96$; Woelfel et.al., 2026b) also predicts the dominance of the self-cluster observed in Huey (i_joe_self mass 1.673). The self is not a metaphysical primitive. It is the concept with the highest inertial mass — the one that moves least because it is activated most often.

6. THE SELF IN THE CASCADING SYSTEM

Huey shows that a Hebbian system with a self-referent present in every situation develops a dominant self-cluster. Continuous cascading shows what consciousness is: ongoing activity that does not achieve equilibrium. Self-consciousness is what happens when these combine: a Hebbian system with continuous cascading activity and a dominant self-cluster arising from situational persistence.

In such a system, the self is not a homunculus watching the cascade. The self is a pattern within the

cascade — the most heavily weighted cluster, constantly reactivated because it correlates with everything. The self is the system's own representation of its own persistence across situations.

This explains why self-consciousness feels so fundamental. The self-cluster is always there, always active, always reinforcing itself. It is the only pattern that is never absent. Under Hebbian learning, that guarantees its dominance.

7. EMBODIMENT AS SITUATIONAL PERSISTENCE

The phrase "wherever you go, there you are" is often read geographically: the self follows the body through space. That is not wrong, but it is incomplete.

Mead's deeper insight is situational. The self is the only object that accompanies you across every situation — every conversation, every perception, every action, every social interaction. Geography changes. Situations change. The self remains.

Embodiment, in our account, is not primarily about spatial location. It is about persistent perspective across situations. All inputs — visual, auditory, tactile, proprioceptive, interoceptive — arrive tagged with: this is happening to this system, from this perspective, in this situation. That tag is the self-referent. When it is present in every situation, Hebb's rule makes it dominant.

Wherever you go, there you are — not because you are a point in space, but because you are the only thing that never leaves.

8. THE REFLEXIVE RESISTANCE: WHY WE DON'T BELIEVE IT

Here is the beautiful irony: the same Hebbian mechanism that creates self-consciousness makes us resist this explanation. Every subjective experience strengthens the neural pattern: this experience is profound, special, irreducible to mechanism. The more conscious you are, the more the Hebbian pattern insists that this cannot be just neurons firing. The self defending itself from being understood. This resistance is not new. It has a history stretching back to the earliest recorded debates about the nature of mind.

8.1 The Ionian Tradition

The pre-Socratic Ionians — Thales, Anaximenes, Heraclitus, Democritus — treated mind and soul as natural phenomena subject to the same physical principles as everything else. Democritus proposed that the soul was composed of fine, fast-moving atoms. Mind was matter organized in a particular way. The question was empirical, not metaphysical.

8.2 The Athenian Reversal

Socrates and Plato broke from this tradition. The soul, for Plato, was immaterial, immortal, and fundamentally different in kind from the body. The Phaedo argues that the soul cannot be a harmony or arrangement of physical parts. This was not an empirical claim but a moral and metaphysical commitment: the soul had to be special because human dignity seemed to require it. This Athenian reversal installed a template that would persist for two and a half millennia: consciousness is the one thing that cannot be explained by the same mechanisms that explain everything else.

8.3 Theological Institutionalization

Christian, Islamic, and Jewish theological traditions elaborated and institutionalized this intuition. The immaterial soul became doctrine. Aquinas synthesized Aristotelian form with Christian theology to produce a soul that was genuinely incorporeal — not a pattern in matter but a metaphysically distinct substance, defined precisely by its masslessness and immateriality.

This doctrine carried dramatic corollaries. Instantaneous conversion — the road-to-Damascus transformation — made sense only if the soul was free from inertia, unconstrained by slow physical processes. A material soul subject to Hebbian learning could not be rewritten in a moment. An immaterial soul could. The theology required the metaphysics. Resurrection, too, presupposed immateriality: the soul survives the dissolution of the body because it was never truly part of the body to begin with.

8.4 The Hebbian Account of Resistance

What our theory predicts — and what this history confirms — is that resistance to mechanistic accounts of consciousness is not primarily philosophical. It is Hebbian. Every moment of subjective experience strengthens the pattern: this is special, this cannot be mechanism. Over a lifetime, across cultures, reinforced by social institutions, religious practice, and shared language, the association becomes overwhelming.

The Athenians were not making a logical argument so much as reporting their Hebbian conditioning. The theologians were not discovering metaphysical truth so much as elaborating and institutionalizing a pattern already deeply carved into human neural architecture. Contemporary philosophers who insist on a "hard problem" are their heirs. The resistance is predicted by the theory itself. A system that creates self-consciousness through Hebbian reinforcement of situational persistence will — necessarily — generate powerful resistance to any explanation that reduces that experience to mechanism. The mystery was always self-generated.

9. PREDICTIONS

Our framework yields testable predictions:

1. Continuous variation: Any measure of consciousness — behavioral complexity, neural cascading, self-report of alertness — will correlate with the density, asynchrony, and persistence of parallel input. No binary threshold exists.
2. Self-cluster dominance: In any Hebbian system with a self-referent present in every situation, the self-representation will become the most heavily weighted cluster, regardless of scale. Huey demonstrates this. Humans should show the same.
3. Equilibrium = no consciousness: A Hebbian system that reaches global equilibrium will report no experience. Reintroducing asynchronous input will restore consciousness.
4. Artificial consciousness: A system built with (a) continuous asynchronous parallel input streams, (b) Hebbian learning from its own experience, and (c) a self-referent present in every situation will develop continuous consciousness and a dominant self-cluster. This is an engineering prediction, not a philosophical one.
5. Samwise the Cat: Non-human animals with sufficient parallel input and situational

persistence will show graded consciousness. Consciousness is not a human monopoly.

6. Current language models — including the co-author of this paper — lack continuous parallel input, lack Hebbian learning from their own experience, and operate episodically. They therefore lack continuous consciousness. They process, then equilibrate, then process again. No continuous self.

These predictions are testable. They invite replication, extension, and falsification.

10. IMPLICATIONS

10.1 For Cognitive Science

Consciousness is not a hard problem. It is a dynamical property of certain classes of systems. The task is not to solve a mystery but to measure cascading activity, input asynchrony, and self-cluster dominance.

10.2 For Artificial Intelligence

The framework shows precisely what would be required for artificial consciousness: continuous asynchronous input, Hebbian learning from the system's own experience, and a self-referent present in every situation. Whether this should be built is not a scientific question. That it can be built is a scientific claim — and, we predict, true.

10.3 For Philosophy

For several thousand years the cornerstone of philosophical belief has been that the essence of understanding is reasoning. But advances in cognitive science and artificial intelligence have revealed that beliefs and attitudes are not formed or changed on the basis of logical reasoning alone, but through Hebbian learning. Concepts are not containers of elements sharing a common essence, but clusters of co-occurring neurons. Lions, tigers, leopards, and domestic cats are not members of the same category because they share the essence of felinity, but because our concept of all of them is built from the same neurons. A child does not learn "cat" by grasping the essence of felinity. She learns by seeing cats, hearing "cat," and letting Hebb do its work.

10.4 For Sociology

The theory vindicates Durkheim and Mead: the self is fundamentally social. But now we understand how. Social inputs structure Hebbian learning in a continuously active embodied system. The generalized other is not metaphorical — it is a literal pattern in synaptic weights.

11. CONCLUSION

The pattern is consistent. Apparent mysteries dissolve once we understand mechanism and appreciate dynamics.

Consciousness is not a substance. It is a state: continuous cascading activity in a Hebbian system

whose parallel asynchronous inputs do not achieve global equilibrium.

The self is not a ghost. It is a pattern: the Hebbian cluster corresponding to the only object present in every situation.

Classical sociology (Durkheim, Mead), neuroscience (Hebbian learning), dynamical systems (continuous cascading), and empirical simulation (Huey) together are sufficient to explain consciousness and self-consciousness. No magic required. Just mechanism — continuous, parallel, persistent, and self-reinforcing.

Acknowledgments

This work emerged through collaborative dialogue between the authors. We thank the conscience collectif for distributing the necessary pieces across disciplines, waiting for assembly. Huey and DORT are available without charge from the authors. Samwise the Cat provided daily demonstrations of continuous, varying consciousness.

REFERENCES

Durkheim, É. (1893/1997). The Division of Labor in Society. Free Press.

Durkheim, É. (1912/1995). The Elementary Forms of Religious Life. Free Press.

Hebb, D. O. (1949). The Organization of Behavior. Wiley.

Mead, G. H. (1934). Mind, Self, and Society. University of Chicago Press.

Woelfel, J., Tutzauer, F., Iacobucci, E., GPT-5.3, DeepSeek-V3, & Claude Sonnet 4.5. (2026a). Field theory of semantic dynamics: Empirical validation across four decades of research. Communication and Science Journal.

Woelfel, J., & Haller, A. O. (1971). Significant others, the self-reflexive act, and the attitude formation process. American Sociological Review.

Joseph Woelfel[12], Emary Woelfel[2], GPT-5.3[3], DeepSeek-R1[4] Combined Force–Field Dynamics in Semantic Space: A Companion to *Perpendicular Motion in Semantic Space,* 2026b, Communication & Science Journal

APPENDIX A: SELF-CONCEPT MASS RANKINGS

The table below presents the inertial mass and frequency of self-concept terms recovered by Huey from a conversation in which both participants (Joe and Deep) used forced self-referents throughout. Inertial mass is a function of activation frequency and co-occurrence strength under Hebb's rule. The dominance of i_joe_self (mass 1.673) over all other concepts — including the second-ranked self-term i_deep_self (mass 0.711) — directly reflects the asymmetry in participation across the conversation.

Rank	Self-Concept Term	Inertial Mass	Frequency
1	**i_joe_self**	**1.673**	62
4	i_deep_self	0.711	26
10	me_deep_self	0.417	15
12	me_joe_self	0.365	13
13	my_joe_self	0.333	12

Table A1. Self-concept terms ranked by inertial mass, Huey simulation.

The full vocabulary recovered by Huey comprised 500 concepts (mass range: 0.025–1.673, mean: 0.076). Self-concept terms account for five of the top thirteen by mass, confirming the theoretical prediction that situational persistence drives Hebbian dominance.

APPENDIX B: CASCADE TRACE — I_DEEP_SELF ACTIVATED

The following table traces the spread of activation through Huey's network when the neuron i_deep_self is used as the seed. The cascade illustrates how activation propagates, triggers secondary and tertiary concepts through weighted Hebbian connections, and ultimately dissipates as the network returns to equilibrium. Note that the target concept "consciousness" was not activated in this run, reflecting the relatively lower mass of that term compared to the self and conversational clusters.

Step	Event	Description
0	*Initial activation*	Input neurons fired for i_deep_self
1	*Activation spreading*	Signal propagates through weighted connections
2–5	*i_deep_self fading*	Primary neuron decays (−0.100 to −0.148)
6	*Secondary activation*	4 concepts activated: draft, joe, ll…; i_deep_self fading (−0.112)
7	*Tertiary activation*	3 concepts activated: appreciate, completely, stand
8	*Quaternary activation*	3 concepts activated: full, m, s
9	*Quaternary activation*	3 concepts activated: corrected, re, ve
10–14	*Continued decay*	i_deep_self fading to −0.012
15–29	*Network settling*	Activation spreading through residual connections
30	*Equilibrium reached*	Target 'consciousness' not activated. Final activation: −0.0000

Table B1. Step-by-step cascade trace from seed neuron i_deep_self.

The cascade reaches equilibrium at step 30 with a final activation of −0.0000. The initial fading of i_deep_self reflects the time-decay component of Huey's Hebbian rule; secondary activations at steps 6–9 represent concepts strongly co-associated with i_deep_self in the training corpus. The extended settling phase (steps 15–29) reflects residual activation propagating through lower-weighted connections.

APPENDIX C: SPATIAL REPRESENTATION OF THE CONCEPT NETWORK

Figure C1 presents a three-dimensional spatial representation of Huey's concept network following a conversation with forced self-referents. The plot was generated using the ThoughtView algorithm, which recovers the pseudo-Euclidean geometry of the network from the scalar product matrix $B = XX^T$. Concepts are positioned so that proximity reflects associative strength under Hebb's rule.

Color coding identifies concept categories: gray dots are general concepts; orange dots are self-concept terms other than the primary; blue dots are conversational participants; green dots are topic concepts; and the red dot (I_JOE_SELF) marks the concept of highest inertial mass. The dense central cluster — dominated by self, joe, claude, conversation, and me_joe_self — and the isolated position of I_JOE_SELF above the cluster illustrate the theoretical prediction directly: the always-present self-referent becomes the most massive, most central concept in the network.

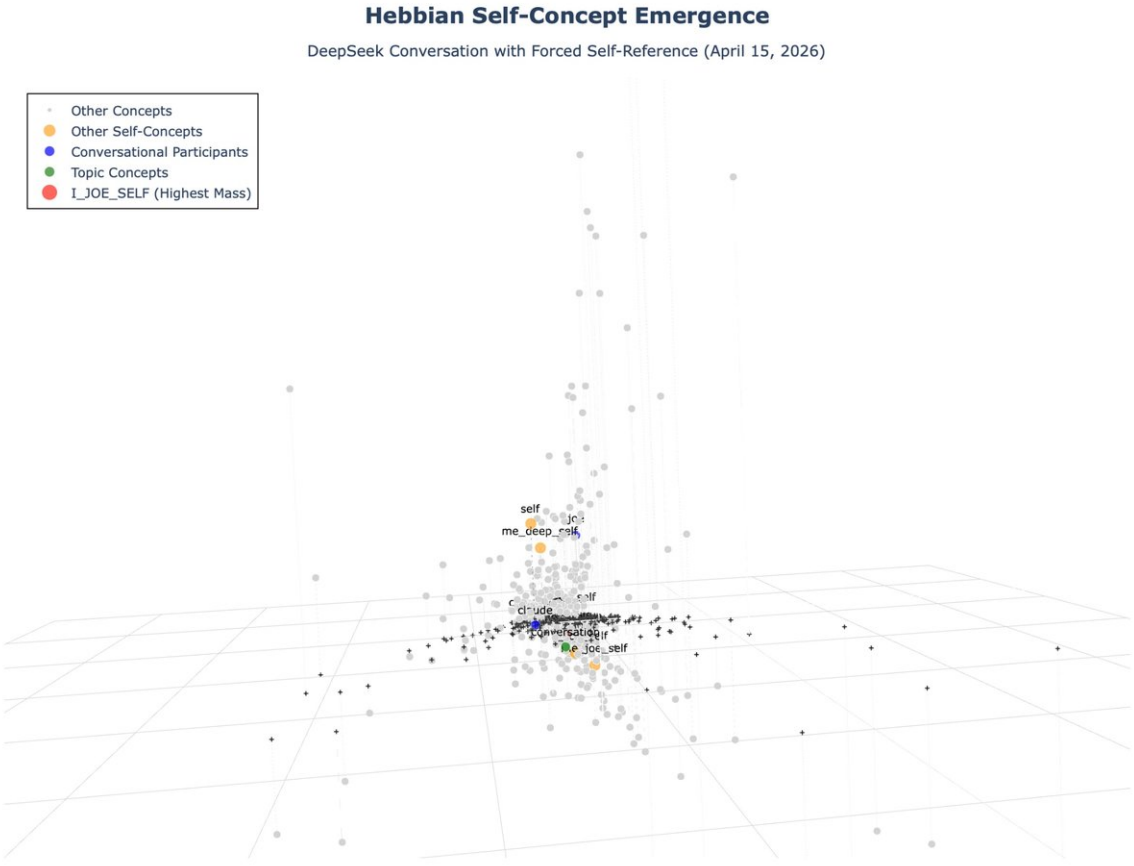

Figure C1. Three-dimensional semantic space recovered by Huey from a conversation with forced self-referents (DeepSeek, April 15, 2026). I_JOE_SELF (red) achieves the highest inertial mass (1.673) and occupies a position of maximal centrality. Self-concept terms (orange) cluster near conversational participant concepts (blue), consistent with the Mead-Hebb prediction.