

From Neuromaps to Meaning: Contextual Duration, Feedforward Biocognition, and the Limits of Artificial Intelligence

Mario Martinez, PsyD
Biocognitive Science Institute
April 2026
bsi@biocognitive.com

Abstract

Contemporary attempts to align artificial intelligence with human cognition have largely proceeded through bottom-up strategies that seek to reconstruct intelligence by scaling upward from neural components. While such approaches have yielded significant advances in perception, classification, and pattern recognition, they encounter a fundamental epistemological limitation when confronted with the problem of meaning. Meaning is not an emergent property of neural activity alone, nor can it be inferred from neuromaps, however sophisticated. This paper advances a top-down, world-relative framework for understanding meaning, arguing that meaning must be established at the level of lived experience and contextual duration before neural or computational explanations can be meaningfully interpreted. Neural systems are reframed as feedforward, anticipatory structures that create the potential for meaning in unknown future contexts rather than as repositories of meaning itself.

The paper first examines the epistemological limits of bottom-up approaches, then develops a top-down framework grounded in contextual duration, and finally considers the implications of this model for artificial intelligence.

Contextual duration is approached not as chronological time but as a temporally extended organization of significance, reflected in anticipatory, context-sensitive engagement with uncertain environments.

The implications of this framework for neuroscience and artificial intelligence are explored, with particular emphasis on the limits of biological mimicry and the necessity of contextual constraint.

Keywords: meaning-making, biocognitive epistemology, neuromaps, artificial intelligence

Introduction

Artificial intelligence research has increasingly turned to neuroscience for guidance, attempting to replicate or approximate human intelligence by modeling neural architectures, cortical hierarchies, and representational pathways (Corrado et al., 2018; Friston, 2010). These efforts are often motivated by the assumption that meaning, cognition, and intelligence can be reconstructed by sufficiently detailed mappings of neural activity. Yet despite impressive technical progress, such approaches remain unable to account for the phenomenon that is central to human intelligence: meaning.

The limitation is not merely computational. It is epistemological. Meaning does not reside in neurons, nor does it emerge automatically from complexity. Functional neuroimaging can correlate neural activity with the occurrence of thought, but it cannot capture the meaning-making of that thought (Pessoa, 2017). Meaning is constituted through context, history, embodiment, and lived duration (Bergson, 1911; Gallagher & Zahavi, 2012). Neural activity supports cognition, but meaning arises only when cognition is situated within a world context (Noë, 2009).

This paper proceeds by first examining the epistemological limitations of bottom-up approaches to cognition, then advancing a top-down, world-relative framework grounded in contextual duration, and finally considering the implications of this framework for artificial intelligence.

The Epistemological Ceiling of Bottom-Up Approaches

Bottom-up models assume that as complexity increases, context accumulates. This assumption fails at precisely the point where meaning appears. Neurons operate in constrained biochemical

environments governed by electrochemical dynamics (Buzsáki, 2006). Cognition, by contrast, unfolds in open environments saturated with social norms, ethical stakes, symbolic inheritance, and temporal irreversibility (Heidegger, 1962). The transition between these levels is not additive. Context is introduced at higher levels of organization; it is not assembled from lower ones.

As one ascends from neural activity to lived experience, the complexity of contextual relevance expands exponentially. Conversely, as one descends from meaning toward neural explanation, contributory contextual complexity is necessarily lost. This loss is structural, not accidental. Neural data can describe support conditions for cognition, but it cannot encode normativity, value, or significance (Anderson, 2014). Two individuals may exhibit similar neural activation patterns while the meaning of their thoughts diverges radically due to differences in biography, intention, or consequence.

Although some hybrid artificial intelligence models combine bottom-up learning strategies with top-down symbolic or rule-based frameworks, they remain constrained by formally specified domains and do not possess the lived contextual duration through which meaning is enacted in human cognition (Tse, 2013).

Top-Down Approaches and World-Relative Meaning

To move beyond these limitations, a different explanatory starting point is required—one that begins with meaning as it is lived rather than as it is inferred from underlying mechanisms.

A top-down approach begins where meaning appears: in lived experience. From this vantage, neural investigation becomes explanatory rather than constitutive. Once meaning is established at the level of context, neural correlates can be identified without mistaking correlation for origin. The error of many contemporary approaches lies not in studying neural mechanisms, but in assigning them generative authority over meaning.

Meaning is world-relative before it is brain-relative (Martinez, 2001). It arises in relation to uncertainty, vulnerability, consequence, and duration (Bergson, 1911). Brains evolved as solutions to contextual problems, not as containers of meaning. Any account of intelligence that neglects this ordering risks confusing infrastructure with phenomenon.

Although contextual duration is not directly measurable as a discrete variable, it can be approached indirectly through its functional signatures. At the human level, contextual duration may be inferred wherever present interpretation reflects an integration of remembered past, lived present, and anticipated future consequence rather than immediate stimulus-response coupling. Such integration becomes evident in narrative continuity, delayed decision-making under uncertainty, context-sensitive modulation of perception, and the capacity to preserve ambiguity until consequence clarifies relevance. In this sense, contextual duration is not chronological time but temporally extended significance. Its empirical trace lies not in duration as quantity, but in the manner in which meaning organizes perception, valuation, and response across unfolding contexts.

Neuromaps as Feedforward Precursors

Within this top-down framework, neural architectures can be reinterpreted not as sources of meaning but as anticipatory systems that make meaning possible.

Reframed within this top-down perspective, neuromaps assume a different role. Rather than encoding meaning, neural architectures create feedforward conditions that enable meaning to be enacted in the future. The brain is an anticipatory system optimized for openness, adaptability, and preparedness in the face of unknown contexts (Tse, 2013).

At the neural level, what exists are precursors of meaning—biocognitive readiness structures that expand degrees of freedom rather than specify semantic content (Tse, 2016). These feedforward

systems sacrifice specificity for flexibility. They do not predict meaning; they prepare the organism for meaningful encounter in uncertain future contexts (Martinez, 2016).

Implications for Artificial Intelligence

Artificial intelligence systems, like neural systems, operate within structured domains. Unlike humans, however, they do not inhabit worlds. They lack embodiment, temporal finitude, ethical consequence, and lived history (Noë, 2009). As a result, artificial systems cannot generate meaning, even when their outputs appear semantically rich.

This argument differs from embodied and enactive approaches that have already challenged purely computational accounts of cognition by emphasizing sensorimotor coupling, world-involvement, and the dependence of intelligence on environmental interaction. Those approaches correctly reject the view that cognition can be exhaustively modeled as internal representation. However, the present framework advances a further claim: embodiment and interaction alone do not yet account for meaning unless they are situated within contextual duration. A system may be dynamically coupled to an environment and still lack the temporally extended horizon through which vulnerability, consequence, ambiguity, and anticipation acquire lived significance. The decisive issue is not only whether cognition is embodied or enacted, but whether it unfolds within a durational field in which relevance is shaped in advance by feedforward openness to uncertain futures.

One consequence of this distinction is that contemporary artificial intelligence systems can exhibit linguistic fluency, pattern sensitivity, and even apparent insight without engaging in meaning-making as humans do. Large-scale models may generate contextually appropriate responses by exploiting statistical regularities across vast corpora, yet such outputs remain detached from lived consequence, vulnerability, and temporal commitment. The appearance of understanding arises

from successful construction within formal domains, not from participation in a world. Without embodied stakes or durational exposure to uncertainty, artificial systems cannot distinguish relevance from irrelevance in the human sense, nor can they recognize when ambiguity should be preserved rather than resolved. This limitation underscores why refinement of human instruction, contextual framing, and interpretive constraint is more consequential for meaningful AI use than further attempts at biological imitation.

Attempts to improve AI by replicating neural architectures misunderstand the locus of meaning and overestimate the relevance of biological mimicry (Lake et al., 2017). Progress toward human-relevant artificial intelligence depends less on bottom-up replication and more on top-down constraint. Artificial systems need not possess consciousness or experience to participate meaningfully in human domains, but they must be guided by principles that preserve contextual relevance. This requires restraint, stopping rules, and sensitivity to ambiguity rather than exhaustive resolution (Vaihinger, 1911; Salk, 1983).

Conclusion

Meaning cannot be reverse-engineered from parts. It must be encountered where it appears: in the world, across duration, through embodied participation.

Contextual duration, in this sense, is not reducible to chronological time but refers to the temporally extended organization of significance through which relevance is sustained, revised, and enacted across unfolding contexts.

Neural systems prepare organisms for that encounter by creating feedforward conditions of openness, adaptability, and anticipatory readiness. They do not, and cannot, contain meaning itself. Neuromaps describe enabling conditions for cognition, but meaning emerges only when those conditions are enacted within concrete contexts shaped by history, consequence, and uncertainty.

Artificial intelligence exposes this distinction with particular clarity. The inability of artificial systems to generate meaning is not a technical deficit to be remedied by increased computational power, larger datasets, or closer emulation of neural architectures. Rather, it reflects a categorical difference between systems that inhabit lived worlds and systems that operate within formally specified domains. Attempts to replicate contextual brain-based learning through algorithmic, linear construction misconstrue both the nature of human meaning-making and the capacities of artificial systems.

It is proposed in this paper that, given the fundamental differences in learning between artificial intelligence systems and the human brain, efforts to make AI “think like the brain” are less productive than approaches that focus on how humans communicate with and guide algorithmic intelligence. Contextual learning in the brain is not a process that can be replicated by computational architectures. However, aspects of human meaning-making can be mimicked when artificial systems are constrained, instructed, and directed through carefully refined conceptual and linguistic frameworks.

From this perspective, progress in human–AI collaboration depends less on imitation and more on translation. Language becomes the critical interface—not as a medium for encoding meaning into machines, but as a means by which humans can impose contextual relevance, interpretive boundaries, and purpose on algorithmic processes. Such an approach preserves the distinctive strengths of artificial intelligence while avoiding the category error of attributing to it capacities that belong to embodied, temporally finite organisms.

References

Anderson, M. L. (2014). *After phrenology: Neural reuse and the interactive brain*. MIT Press.

Bergson, H. (1911). *Time and free will: An essay on the immediate data of consciousness* (F. L. Pogson, Trans.). George Allen & Unwin. (Original work published 1889)

Bergson, H. (2007). *Creative evolution* (A. Mitchell, Trans.). Dover Publications. (Original work published 1911)

Buzsáki, G. (2006). *Rhythms of the brain*. Oxford University Press.

Clark, A. (2016). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.

Corrado, G. S., Ling, L., Matheson, T., & others. (2018). The dawn of deep learning for artificial intelligence. *IEEE Computer*, 51(9), 44–52. <https://doi.org/10.1109/MC.2018.3620972>

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138. <https://doi.org/10.1038/nrn2787>

Gallagher, S., & Zahavi, D. (2012). *The phenomenological mind* (2nd ed.). Routledge.

Heidegger, M. (1962). *Being and time* (J. Macquarrie & E. Robinson, Trans.). Harper & Row. (Original work published 1927)

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, e253. <https://doi.org/10.1017/S0140525X16001837>

Martinez, M. E. (2001). The process of knowing: A biocognitive epistemology. *Journal of Mind and Behavior*, 22 (4), 407–426.

Martinez, M. E. (2016). *The MindBody Code: How to change the beliefs that limit your health, longevity, and success*. Sounds True.

Noë, A. (2009). *Out of our heads: Why you are not your brain, and other lessons from the biology of consciousness*. Hill and Wang.

Pessoa, L. (2017). *The cognitive-emotional brain: From interactions to integration*. MIT Press.

Salk, J. (1983). *Anatomy of reality: Merging of intuition and reason*. Columbia University Press.

Tse, P. U. (2013). *The neural basis of free will: Criterial causation*. MIT Press.

Tse, P. U. (2016). Criteria-based causation without reductionism. *Progress in Brain Research*, 228, 197–219. <https://doi.org/10.1016/bs.pbr.2016.03.010>

Vaihinger, H. (1911). *The philosophy of “as if”* (C. K. Ogden, Trans.). Routledge & Kegan Paul.