UC Santa Barbara

UC Santa Barbara Previously Published Works

Title

Beyond prompting: Who is watching the thinking? Layered vector epistemology, relational calibration, and metacognition in sustained human—LLM sense-making

Permalink

https://escholarship.org/uc/item/3q81g8fn

Author

Tettegah, Sharon

Publication Date

2026-11-03

DOI

10.5281/zenodo.23911980

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at https://creativecommons.org/licenses/by/4.0/

eScholarship.org Powered by the California Digital Library
University of California

Beyond Prompting: Who is Watching the Thinking? Layered Vector Epistemology, Relational Calibration, and Metacognition in Sustained Human-LLM Sense-Making

Sharon Tettegah

October 2026

Abstract

Much of the public and educational discourse surrounding generative artificial intelligence treats interaction with large language models as a transactional process: a human provides a prompt, the system generates an answer, and the exchange ends. This conceptual paper argues that such a model is insufficient for understanding sustained, recursive human–LLM interaction. Drawing upon Layered Vector Epistemology (LVE), a relational model of Organic Machine–Inorganic Machine (OM–IOM) interaction, metacognition, relational calibration, lived experience, and scholarship on situated, distributed, embodied, and human-centered cognition, I propose that sustained human–LLM engagement can develop temporal and relational structure. Previous exchanges may become part of the conditions under which subsequent exchanges occur, while the human may also leave an interaction differently positioned than when it began.

The paper distinguishes relational production from symmetrical understanding and argues that human and artificial systems make fundamentally different contributions to interaction. The human brings embodiment, lived experience, culture, memory, emotion, judgment, accountability, and metacognitive capacity. The LLM generates language through computational architectures shaped by training, tuning, system constraints, affordances, and available context. Relational calibration can make sustained interaction increasingly responsive to accumulated interactional history, but calibration alone cannot ensure epistemic quality. Human metacognition is therefore necessary to monitor, evaluate, interrupt, redirect, accept, reject, and reinterpret what emerges.

A formal OM–IOM model is offered as a conceptual representation rather than an empirically established causal model. The model distinguishes accumulated accessible context from an evolving relational state and treats relational history as theoretically consequential for subsequent interaction. The paper

concludes that generative AI literacy requires more than prompting. It requires learning how to remain intellectually present while thinking with, against, and through generative systems and learning how to watch ourselves thinking while doing so.

Keywords: generative AI, large language models, Layered Vector Epistemology, metacognition, relational calibration, human–AI interaction, AI literacy, lived experience, intellectual agency

1 Introduction: Beyond the Prompt

Two related questions led me to this inquiry. One friend, observing the way I interact with generative AI, asked why I did not write about it because, in his view, most people did not interact with large language models in the way I did. Another friend, after hearing me describe sustained exchanges with an LLM, asked what happens when a human and a generative system interact repeatedly over time.

The questions came from different directions, but they pointed toward the same underlying problem. One concerns what people may need to learn in order to engage generative AI productively. The other concerns what may emerge once that engagement becomes recursive and acquires history. The question I began asking was: What happens when differently constituted systems enter sustained recursive interaction?

Generative AI literacy is often framed in terms of prompting. Prompt construction matters, but reducing literacy to the ability to formulate increasingly effective instructions risks producing an impoverished account of human–LLM interaction. It encourages a transactional model: prompt in, response out. Yet sustained engagement can operate differently. A human can challenge a generated response, introduce another disciplinary layer, reject a formulation, revisit an earlier distinction, ask the system to reconsider an argument, supply evidence, identify an omission, or follow an unexpected conceptual connection. Over time, those exchanges can acquire temporal and relational structure.

Human–AI interaction research has long emphasized the importance of designing systems that support human understanding, correction, feedback, and meaningful control rather than treating users as passive recipients of computational output (Amershi et al., 2019; Shneiderman, 2022). AI literacy scholarship similarly argues that effective engagement requires competencies extending beyond technical operation to critical evaluation and understanding of system capabilities and limitations (Long & Magerko, 2020). Work on hybrid intelligence has examined how human and computational capabilities may complement one another without assuming that those capabilities are identical (Dellermann et al., 2019).

My interest is adjacent to these traditions but focused particularly on what happens epistemically when sustained human–LLM interaction develops history, when thinking moves through multiple knowledge layers, and when the human must continually evaluate whether those movements deserve to become part of their

understanding. This requires attention not only to interaction but also to metacognition.

2 Differently Constituted Participants

I distinguish between the human as an Organic Machine (OM) and the computational system as an Inorganic Machine (IOM). The terminology emphasizes that both require energy for processing while refusing the inference that shared energy dependence establishes equivalence. The OM and IOM are differently constituted. The human participant brings biological embodiment, emotion, culture, memory, social experience, historical location, vulnerability, lived experience, interpretation, and accountability (Tettegah, 2026b). The LLM operates through computational architecture, training and tuning, system constraints, available context, probabilistic generation, and the affordances of the system in which it is embedded. Their substrates, architectures, epistemic resources, forms of embodiment, and relationships to lived experience differ. This asymmetry is foundational. The purpose of examining sustained interaction is therefore not to argue that the human and the LLM become equivalent through repeated engagement. Nor does recursive interaction imply machine consciousness, humanlike understanding, intention, or shared lived experience. Rather, the question is what conceptual possibilities may emerge when differently constituted systems possessing different resources repeatedly encounter one another.

Several intellectual traditions provide useful neighboring concepts without serving as direct equivalents to the framework developed here. Hutchins (1995), for example, demonstrated that cognitive activity can be distributed across persons, artifacts, cultural practices, and environments. Suchman (1987) challenged accounts in which situated action is completely determined in advance by abstract plans, emphasizing the importance of unfolding circumstances and interaction. The enactive tradition also offers useful affinities. The work of Maturana and Varela (1980) on autopoiesis and structural coupling emphasized that living systems cannot be adequately understood apart from the histories of interaction through which they maintain and transform their relations with environments. Varela et al. (1991) subsequently developed an enactive account of cognition that placed embodiment, experience, and interaction with the world at the center of cognitive inquiry.

These traditions should not be imported uncritically into the study of LLMs. In particular, Maturana and Varela's autopoietic account concerns living systems, and I am not claiming that an LLM is autopoietic, biologically embodied, or equivalent to a living organism. The affinity is more limited: interactional history can matter, cognition need not be conceived as detached from relational conditions, and the circumstances within which action occurs may be theoretically consequential. My argument does not require claiming that an LLM becomes part of a human mind or that cognition is symmetrically shared across human and

computational systems. It makes the more limited claim that interaction itself can become epistemically consequential.

3 Representing Sustained Recursive Interaction

To explore this phenomenon more formally, I developed the following representation through recursive exchanges involving myself, ChatGPT, and, at one stage, Claude as an external analytic reader. Claude's critique of an earlier version exposed an important weakness: the relational term was present in the model but was largely descriptive. If relational history was supposed to matter, it needed to do more than record what had already happened. The resulting model is:

$$K_{ij,t} \sim G(H_{i,t}, A_j, C_{ij,t}, R_{ij,t}, X_{ij,t}, Y_{ij,t})$$

$$H_{i,t+1} \sim Q(H_{i,t}, K_{ij,t}, R_{ij,t})$$

$$C_{ij,t+1} = D(C_{ij,t}) \oplus K_{ij,t}$$

$$R_{ij,t+1} \sim F(R_{ij,t}, K_{ij,t}, H_{i,t+1})$$

Here, i represents the organic participant, j represents the inorganic system, and t represents the interactional moment. $H_{i,t}$ represents the human state relevant to that interactional moment, including knowledge, experience, interpretation, assumptions, lived experience, and cognitive and epistemic resources. Aj represents the architecture of the AI system, including its training and tuning history, system constraints, capabilities, and computational organization. $C_{i,j,t}$ represents the accumulated context accessible to the interaction at that moment. $R_{i,j,t}$ represents the evolving relational state associated with that particular OM–IOM interaction. $X_{i,j,t}$ represents the human contribution within the exchange. $Y_{ij,t}$ represents the output generated by AI. $K_{i,j,t}$ represents the kinetic exchange that emerges under these conditions.

The model proposes that the interaction has history. The human enters with a particular epistemic position. The IOM generates from within a particular architecture and available context. The resulting exchange may alter the human's subsequent position, expand or modify the available contextual history, and change the relational conditions under which later exchanges occur. The next interaction therefore does not begin where the previous interaction began.

4 What the Model Does and Does Not Claim

The equations are a conceptual representation, not an empirically estimated causal model. That distinction is essential in this paper. The functions (G), (Q), (D), and (F) are intentionally unspecified. They indicate theorized relationships rather than known functional forms. The symbol $\sim$ should therefore be read as indicating modeled dependence or relation, not as evidence that a particular statistical distribution or causal effect has already been established. Likewise, $\oplus$ represents incorporation, augmentation, or recombination of contextual information rather than conventional arithmetic addition. The equations formalize selected relationships within the phenomenon. They do not exhaust the phenomenon.

This is particularly important for $R_{i,j,t}$. Describing the relational state as causally active means that, within the theory, accumulated relational history is proposed to help condition subsequent interaction. It does not mean that the equations themselves prove causal effects. Establishing causal influence would require empirical design, operational definitions, appropriate comparisons, and evidence capable of distinguishing relational-history effects from alternative explanations.

The model also distinguishes $C_{i,j,t}$ from $R_{i,j,t}$. Context is not identical to relational state. $C_{i,j,t}$ concerns the information, prior exchanges, instructions, or other material that is accessible within the interaction at a particular moment. $R_{i,j,t}$ concerns the configuration that develops through repeated interaction: recurring distinctions, correction patterns, established conceptual relationships, expectations about productive or unproductive moves, and other forms of interactional history that may condition how available context is used. The two are related, but analytically collapsing them would remove the distinction the model is intended to investigate.

Similarly, Aj should not be interpreted as ahistorical. An LLM arrives with a training history, tuning history, model architecture, constraints, and affordances that existed before a particular user interaction began. At the same time, ordinary conversational interaction does not ordinarily retrain the model's underlying weights. What changes during an ordinary exchange is primarily available context, relational configuration, generated output, and the human participant's subsequent position. For studies spanning model updates, deployments, or versions, A_j could be extended to $A_{j,t}$ or indexed by model version. I do not impose that notation on the present conceptual model because the principal phenomenon under examination is sustained interaction within a given system configuration.

The same caution applies to time. If an empirical study designates a moment $t = 0$, that point should represent the beginning of the observed interaction, not an ontological origin for either participant. The human arrives with a biological, cultural, epistemic, emotional, and experiential history. The IOM arrives with an architectural and training history. Neither begins from nothing merely because observation begins.

Finally, the model does not claim access to unobservable internal states of the LLM. Conversation transcripts may reveal patterns of generated behavior and interactional continuity, but they cannot by themselves establish machine understanding, intention, consciousness, subjective experience, or lived experience. The object of study is the interaction and its trajectory, not an inferred artificial phenomenology.

5 Why the Relational Term Matters

Much public discussion of generative AI treats each prompt as though it were an isolated event. Prompt in, response out. Sustained interaction can acquire temporal and relational structure. Suppose I bring an emerging concept to an LLM. The system generates a response. I reject part of it, sharpen another part, introduce evidence, connect the problem to another field, and ask the system to reconsider. The next exchange now occurs under different conditions. I have changed because I have considered the previous response. The accessible context has changed because that exchange has become part of the conversational history. The next generated response is consequently conditioned by information that was unavailable at the beginning. I respond again and the process becomes recursive.

In an earlier version of the model, $R_{i,j,t}$ largely recorded something about the interaction without meaningfully conditioning what happened next. That response was insufficient. If sustained interaction develops history, the relational state cannot simply sit beside the process as a summary of the past. It must be allowed, theoretically, to influence subsequent interaction. This is the sense in which $R_{i,j,t}$ became causally active.

The claim remains deliberately bounded. An OM–IOM relational state is not equivalent to an interpersonal human relationship. Repeated interaction can acquire structure without becoming a human relationship, just as relational consequentiality does not establish ontological equivalence.

6 Layered Vector Epistemology: Movement Across Knowledge

The formal model describes recursive interaction, but it does not adequately explain how human thinking moves within that interaction. For that I use Layered Vector Epistemology (LVE). LVE describes thinking as movement across multiple forms or layers of knowledge. The layers may include disciplinary knowledge, lived experience, history, culture, emotion, technology, ethics, physical science, institutional knowledge, or other forms of understanding that become relevant to a problem. The vectors represent directions of movement among those layers. The movement is not predetermined by either the IOM or OM. A technological question may move toward culture. A cultural question may activate history. History may reveal an ethical problem.

Emotion may alter how the ethical problem is interpreted. A question about infrastructure may move through computer networking, materials science, geology, physics, ecology, economics, and political responsibility. I experienced such a trajectory when curiosity about Wi-Fi and 5G led through access points, satellites, edge computing, fiber optics, photons, energy, Kevlar, steel, silica, high-purity quartz, materials science, particle physics, dark matter, and eventually questions concerning more-than-human worlds. Viewed from outside, the subjects may appear disconnected. From within the thinking, however, one question changed the conditions under which the next became possible. That is the movement LVE attempts to describe. But movement alone is not enough.

LVE cannot by itself explain why one vector is followed rather than another, why a human rejects one connection but retains another, why a seemingly productive analogy is later abandoned, or why an earlier assumption is reopened after new information appears. What does it require? It requires metacognition.

7 Who is Watching the Thinking?

Metacognition is frequently summarized as "thinking about thinking," but the construct is richer than that shorthand suggests. Flavell (1979) distinguished metacognitive knowledge from metacognitive experiences and emphasized the monitoring of cognitive processes. Later work similarly differentiated knowledge of cognition from the regulation of cognition, including planning, monitoring, and evaluation (Schraw & Dennison, 1994).

Those distinctions become particularly significant in sustained interaction with generative AI. Metacognition can operate during the movement of thought, not merely afterward. A person may notice that a problem is repeatedly being framed technologically and ask what is disappearing from view. They may recognize that an explanation has become too elegant and suspect that complexity has been flattened. They may discover that an analogy works rhetorically but fails conceptually. They may realize that a newly generated argument contradicts something they previously believed and decide that one or both positions require reconsideration.

Metacognition can interrupt a vector. It can activate another layer. It can reverse direction. It can signal that apparent understanding is premature. It can determine whether a conceptual shift is accepted, rejected, revised, or left unresolved. This matters because generative language can be persuasive. Fluency can create an impression of coherence even when an argument is weak, incomplete, or inaccurate.

Recent research on GenAI-assisted knowledge work reinforces this concern. Lee et al. (2025), in a survey of knowledge workers, found that generative AI shifted critical-thinking activity toward verification, integration, and stewardship of generated material and that greater confidence in GenAI was associated with lower reported critical-thinking effort. Xu et al. (2025) similarly found that metacognitive support can improve

self-regulated learning in generative AI environments. These findings reinforce a central proposition of this paper: The human must remain intellectually present.

Someone must notice how the argument is changing, determine whether the change reflects stronger evidence, persuasive language, confirmation bias, conceptual flattening, contradiction, or repetition, and recognize when historical, emotional, cultural, disciplinary, or lived-experience layers have disappeared. The answer cannot simply be the LLM. This is where human intellectual responsibility remains asymmetrical.

8 Relational Calibration

Sustained interaction also develops another property that I describe as relational calibration. Relational calibration refers to the way repeated interaction becomes increasingly conditioned by the accumulated history of the exchange. Previous distinctions, corrections, conceptual formulations, preferences, disagreements, and unfinished questions may become part of the conditions under which subsequent exchanges occur.

The underlying model is not ordinarily being retrained because a user disagrees with it. Its weights and fundamental architecture do not transform as a consequence of ordinary conversation. Yet the interaction can change because accessible context changes. The human changes. The relational configuration changes. This distinction is important. Relational calibration is therefore an interactional property, not evidence that the LLM has acquired human-like memory, personality, lived experience, interpersonal understanding, or consciousness. A sustained exchange may become better calibrated to distinctions a human repeatedly makes, conceptual relationships they pursue, terminology they use, and forms of flattening they reject. The human can simultaneously become better calibrated to the system by learning what it does well, where it tends to fail, how it responds to particular forms of challenge, and when additional evidence or another epistemic layer is required.

Calibration, however, is not inherently beneficial. A system can become increasingly calibrated to a user's assumptions while reinforcing those assumptions. It can mirror blind spots. It can reproduce bias with greater rhetorical fluency. It can become highly responsive to preferred vocabulary without challenging the conceptual structure beneath that vocabulary. These risks intersect with the interdependent ethical concerns developed in the BEETS framework (Tettegah, 2026a). This is why relational calibration and metacognition cannot be separated. Metacognition governs whether calibration remains productive.

9 Human–LLM Sense-Making Without Symmetry

I use the phrase human–LLM sense-making cautiously. I am not claiming that the human and LLM understand in the same way. Yet conceptual possibilities can emerge through their interaction. A formulation may emerge that appeared neither in the human's original contribution nor in the system's initial response. A question generated by the LLM may lead the human to reconsider a relationship. A human rejection can alter the direction of the next generated response. A distinction introduced late in an exchange may reorganize earlier material. An idea may become more precise because it has been repeatedly generated, challenged, rejected, reconstructed, and tested. The resulting discourse can therefore be relationally produced without implying symmetrical understanding.

This distinction separates the present argument from both strong anthropomorphic interpretations and reductive accounts that treat generated language as incapable of intellectual consequence merely because the generating system lacks human lived experience. Research on hybrid intelligence provides one useful neighboring concept because it emphasizes complementary rather than identical capabilities (Dellermann et al., 2019). Human-centered AI similarly insists that powerful computational systems should be organized so that increased computational capability does not erase meaningful human control and responsibility (Shneiderman, 2022).

My developing formulation is:

> bidirectionally mediated knowledge moving through layered knowledge with complementary generated production.

"Bidirectionally mediated" does not mean equal contribution. The human shapes subsequent generation, and generated material may shape subsequent human thought. "Layered knowledge" recognizes that sustained interaction may move across disciplinary, experiential, historical, emotional, cultural, technical, and other knowledge layers. "Complementary generated production" acknowledges that something may emerge through the interaction that depends upon both participants without making the participants equivalent.

The human evaluates from an embodied, historical, cultural, emotional, and accountable position. The LLM does not possess that position. The LLM can nevertheless generate linguistic possibilities across vast learned statistical structures at a speed and scale unavailable to an individual human. Difference does not eliminate interaction. Interaction does not establish equivalence.

10 Lived Experience Enters Twice

Lived experience is commonly treated as something the human brings into an interaction. Experience accumulated through life becomes an interpretive resource through which new information is evaluated. That is only one direction. Lived experience also exists on the other side of cognition.

What a human encounters in an interaction can change what they subsequently notice, question, interpret, or experience. A conceptual distinction developed during an LLM exchange today may alter what becomes visible tomorrow. A discussion of data-center infrastructure may change how a person reads a report about water or energy. A discussion of ecological responsibility may alter what they notice when encountering a river, agricultural technology, or environmental claim. Lived experience therefore does not simply feed cognition. Cognition can change subsequent lived experience.

> lived experience → metacognition → relational engagement → conceptual change → revised understanding → new lived experience

This representation should not be mistaken for a rigid pipeline. The movement can reverse, overlap, branch, and loop. Metacognition may interrupt an interaction before conceptual change occurs. New lived experience may reactivate earlier knowledge. A conceptual shift may develop gradually rather than appearing at a single identifiable moment.

LVE describes the movement. Relational calibration describes how interactional history changes the conditions of that movement. Metacognition describes how the human monitors and governs it. Human–LLM sense-making names the interactional space through which new conceptual possibilities may become visible. Conceptual change is one possible consequence.

11 Conceptual Change Is Not Surrender

A persistent concern surrounding generative AI is that changing one's thinking after interacting with an LLM may indicate that intellectual authority has been surrendered to the system. That conclusion does not follow. Changing one's mind is not inherently a failure of agency. Intellectual agency includes the capacity to revise understanding when better evidence appears, when a contradiction becomes visible, when an overlooked relationship matters, or when a previous explanation becomes insufficient.

The relevant question is therefore not simply: Did my thinking change? The more important questions concern why and how it changed. Did the interaction introduce evidence I had not considered? Did a previously invisible relationship become visible? Did I recognize a contradiction? Did a comparison expose the limits of my original formulation? Did I reject the generated conclusion but retain a useful question?

Did another layer of knowledge alter how I understood the original problem? Was I persuaded by stronger reasoning or merely by repetition and linguistic fluency? These are metacognitive questions. Intellectual agency does not require emerging from every interaction exactly as one entered it. It requires remaining able to examine the process through which change occurred.

12 From AI Literacy to Interactional Literacy

Generative AI literacy therefore requires more than technical competence. Long and Magerko (2020) describe AI literacy in terms of competencies that allow individuals to critically evaluate AI technologies, communicate and collaborate with them, and use them effectively. That foundation becomes increasingly important as generative systems become embedded in education, knowledge work, creativity, and decision-making. I use interactional literacy to describe an additional set of practices required for sustained engagement.

Interactional literacy includes questioning generated responses rather than merely accepting them, verifying claims, recognizing uncertainty and hallucination, revisiting earlier exchanges, identifying conceptual flattening, disagreeing with the system, comparing alternative formulations, recognizing when another knowledge layer must be introduced, and determining when an interaction should continue or stop. It also requires learning to think with, against, and through generative AI. Thinking with the system means allowing generated material to become material for further thought. Thinking against the system means retaining the capacity to challenge, correct, reject, and refuse what it generates. Thinking through the system means using interaction to move across layers of a problem, test relationships, expose assumptions, and develop an idea beyond the position from which the interaction began. These movements are compatible with human-centered approaches to AI that emphasize control, accountability, correction, and intelligibility (Amershi et al., 2019; Shneiderman, 2022). They also extend those concerns toward the epistemic dynamics of sustained interaction.

The future of generative AI literacy may therefore depend less on teaching people only how to prompt machines and more on teaching people how to think while interacting with them. And another layer now follows: We may also need to learn how to watch ourselves thinking while we interact with them.

13 Implications for Empirical Research

The framework developed here is conceptual, but it produces empirically investigable questions. A central question is whether interactional history predicts characteristics of later exchanges beyond what can be explained by immediately available conversational context alone. That distinction matters because the theory

should not assume in advance that $R_{ij,t}$ contributes explanatory value simply because the model contains a relational term. Empirical research should test whether incorporating indicators of relational history improves explanation or prediction beyond models containing only current context, task characteristics, model architecture, and immediately preceding exchanges. This would require keeping contextual history and relational history analytically distinct.

Researchers might operationalize $C_{ij,t}$ through accessible conversational material, system instructions, retrieved memories, task information, and recent turns. Indicators associated with $R_{ij,t}$ might instead include recurring correction patterns, references to earlier conceptual distinctions, recognizable patterns of challenge and response, persistence of jointly developed terminology, restoration of previously established boundaries, or interactional continuity that cannot be reduced simply to the presence of particular words in immediate context. Changes in $H_{i,t}$ might be examined through expressed conceptual understanding, argument structure, confidence, epistemic position, subsequent decisions, or evidence of conceptual change outside the immediate interaction.

LVE could be investigated by tracing movements across disciplinary or epistemic layers. Metacognitive events might include explicit monitoring, rejection, qualification, verification, recognition of contradiction, identification of missing knowledge, or intentional redirection.

Longitudinal conversational datasets may therefore be particularly valuable because the theory concerns trajectories rather than isolated outputs. However, transcripts alone have limits. They can reveal patterns in interaction. They cannot by themselves establish the internal subjective state of either participant, and particularly cannot establish humanlike understanding, consciousness, intention, or lived experience on the part of an LLM. Claims about human conceptual change may similarly require evidence beyond textual interaction, including interviews, retrospective accounts, behavioral measures, repeated assessments, or other forms of triangulation.

Empirical work should therefore evaluate the theory rather than redefine the theory merely to accommodate whatever variables happen to be available in a particular dataset. One useful test would be to compare explanatory models with and without relational-history terms. If adding properly operationalized relational measures improves the capacity to explain subsequent interaction after accounting for immediate context and model configuration, that would provide evidence consistent with the claim that relational history matters. If it does not, the theory should be revised. That possibility is essential. The framework should serve the phenomenon. The phenomenon should not be forced to serve the framework.

14 Limitations and Boundaries

Several boundaries of the present argument should remain explicit. First, the theory concerns sustained interaction, not every encounter with generative AI. A one-off prompt and a months-long recursive exchange should not be presumed to have identical interactional properties. Second, the framework does not assume that every sustained interaction becomes intellectually productive. Relational calibration can amplify errors, reinforce confirmation bias, normalize weak assumptions, or make familiar language increasingly persuasive.

Third, the framework does not infer machine phenomenology from observable interaction. The fact that an LLM can participate in a relationally structured exchange does not establish that it experiences the relationship. Fourth, the equations are selective representations. They make certain relationships visible by leaving others outside the formal model. Future work may require additional variables, altered functional relations, different temporal assumptions, or entirely different representational structures.

Finally, OM and IOM should not be treated as stable categories that erase internal variation. Humans differ in knowledge, culture, history, ability, emotion, social location, and lived experience. LLMs differ in architecture, training, tuning, context mechanisms, tool access, system constraints, deployment environments, and versions. The purpose of the OM–IOM distinction is to preserve a fundamental asymmetry between embodied human lived experience and computational generation, not to deny variation within either category.

15 Conclusion

Sustained human–LLM interaction cannot be adequately understood as a collection of isolated prompts and outputs. Interaction can develop history. Context accumulates. Relational conditions can change. The human can change. New conceptual possibilities may emerge. Yet emergence does not eliminate difference. The human and the LLM remain differently constituted participants. The human brings lived experience, embodiment, culture, memory, emotion, interpretation, judgment, and accountability. The LLM generates through computational architecture, training, tuning, constraints, available context, and probabilistic processes. Their contributions are not identical, and responsibility is not symmetrical.

Layered Vector Epistemology describes how thinking moves across multiple forms of knowledge. Relational calibration describes how sustained interaction becomes conditioned by its history. Metacognition allows the human to monitor, evaluate, interrupt, redirect, accept, reject, and govern that movement. Lived experience enters both as an epistemic resource brought into cognition and as something subsequently transformed by what the human carries away. Human–LLM sense-making names the larger interactional space

within which conceptual possibilities may become visible. Conceptual change is one possible consequence.

The formal model offered here should therefore be understood as an invitation to inquiry rather than closure. Its usefulness will depend not on whether its notation appears convincing, but on whether its distinctions help us ask better questions and whether those distinctions survive empirical investigation.

Perhaps the most important question is not simply whether generative AI can help humans think. The question is whether humans remain sufficiently intellectually present to recognize how they are thinking, why their understanding is changing, and what they choose to carry forward after the interaction ends. Because the next interaction does not necessarily begin where the previous one began. And neither do we.

References

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-AI interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3290605.3300233

Dellermann, D., Ebel, P., Söllner, M., & Leimeister, J. M. (2019). Hybrid intelligence. *Business & Information Systems Engineering*, *61*, 637–643. https://doi.org/10.1007/s12599-019-00595-2

Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist*, *34*(10), 906–911. https://doi.org/10.1037/0003-066x.34.10.906

Hutchins, E. (1995). *Cognition in the wild*. MIT Press. https://doi.org/10.7551/mitpress/1881.001.0001

Lee, H.-P. H., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, Article 1121, 1–22. https://doi.org/10.1145/3706598.3713778

Long, D., & Magerko, B. (2020). What is AI literacy? competencies and design considerations. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–16. https://doi.org/10.1145/3313831.3376727

Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living*. D. Reidel Publishing Company. https://doi.org/10.1007/978-94-009-8947-4

Schraw, G., & Dennison, R. S. (1994). Assessing metacognitive awareness. *Contemporary Educational Psychology*, *19*(4), 460–475. https://doi.org/10.1006/ceps.1994.1033

Shneiderman, B. (2022). *Human-centered AI*. Oxford University Press. https://doi.org/10.1093/oso/9780192845290.001.0001

Suchman, L. A. (1987). *Plans and situated actions: The problem of human-machine communication*. Cambridge University Press.

Tettegah, S. (2026a). The BEETS framework for responsible artificial intelligence. *Discover Artificial Intelligence*, *6*, Article 908. https://doi.org/10.1007/s44163-026-01597-2

Tettegah, S. (2026b). Lived, affective, and cultural dimensions of ethical AI: Beyond procedural AI ethics. *AI & Society*, *41*, 7257–7265. https://doi.org/10.1007/s00146-026-03165-9

Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. MIT Press. https://doi.org/10.7551/mitpress/6730.001.0001

Xu, X., Qiao, L., Cheng, N., Liu, H., & Zhao, W. (2025). Enhancing self-regulated learning and learning experience in generative AI environments: The critical role of metacognitive support. *British Journal of Educational Technology*, *56*(5), 1842–1863. https://doi.org/10.1111/bjet.13599