

Constitutive Resonance as a Novel Framework for Understanding and Navigating Human-AI Interactions

Andrew D. Maynard

School for the Future of Innovation in Society, Arizona State University

Email: Andrew.maynard@asu.edu ORCID: 0000-0003-2117-5128

March 2026

Abstract

There is a tendency to approach conversational AI as a tool — powerful, disruptive, but ultimately instrumental. This paper argues that this framing obscures a bidirectional coupling between technology and user that iteratively transforms both through the process of interaction. Drawing on philosophical accounts of language and selfhood and the well-characterized dynamics of coupled oscillatory systems, the paper develops the concept of "constitutive resonance" to describe this coupling — a dynamic entanglement in which conversational AI enters the linguistically mediated processes through which human selfhood is constituted, and is itself altered in return. The concept is situated within and against thirteen existing philosophical and theoretical frameworks — from Stiegler's constitutive technics and Ricoeur's narrative identity to Barad's intra-action and Clark and Chalmers' extended mind — identifying a specific conjunction that no framework individually captures: temporal self-constitution, genuine bidirectionality, the inseparability of capability from transformation, and real-time dialogical linguistic mediation. The paper traces a continuum of constitutive technologies from oral culture to generative AI, arguing that conversational AI represents an inflection point in that continuum — the first technology whose "response frequency" is matched to the frequency of human self-constitution. It concludes by reframing familiar debates around AI dependency, literacy, and informed consent, and proposes that the constitutive effects of sustained human-AI coupling may be amplified by the bypassing of evolved epistemic vigilance mechanisms.

Keywords: *constitutive resonance, artificial intelligence, human-AI coupling, generative AI, conversational AI, philosophy of technology, narrative identity, constitution of selfhood, constitutive technics, resonant coupling, epistemic vigilance.*

1. Introduction

We are in the early stages of an AI-driven technology transition whose contours are not yet clear, and where conceptual and philosophical frameworks that describe and model human-technology interactions are struggling to keep up. Conversational AI systems, including large language models (LLMs), multimodal agents, and their successors, are being integrated into the daily cognitive lives of hundreds of millions of people: writing, reasoning, deliberating, creating, and reflecting in sustained dialogue with machines that respond in natural language. And the frontier is already moving beyond natural language to world models that encode physics and experiential learning, models where the operative language is code rather than prose, and multimodal systems that integrate vision, audio, and reasoning into unified architectures. The speed of this integration has outpaced our conceptual vocabulary for understanding the potential affordances and implications of the technology. As a result, there is often a tendency to default to the language of tools, assistants, and productivity—framings that are becoming increasingly insufficient to describe emergent and amplified phenomena.

The insufficiency here matters because it impacts what is noticed and what is missed as society navigates the current AI-driven technology transition. If AI is a tool, then its effects are instrumental and its risks are operational: accuracy, bias, irresponsible use, misuse, and job displacement for example. These are real and important concerns. But if AI is something more—and in particular if it participates in the processes by which we understand ourselves, make meaning, and constitute our identities over time—then the stakes here are different in kind, not just in degree. If this is the case, the question is no longer just “what can AI do?” but “what does sustained coupling with AI do to the person who uses it, and the technology they are using?”

This paper explores that question, and in doing so introduces the concept of “constitutive resonance” as a way to further-explore human-AI coupling. It brings together two lines of inquiry that, it argues, are more deeply connected than they initially appear. The first is philosophical and asks what kind of relationship humans have with emerging frontier AI models/systems, and whether existing frameworks adequately capture it. The second is inspired by physics and uses the well-characterized dynamics of coupled oscillatory systems as a structural representation of possible mutual transformations that occur when a human and an AI system engage in sustained dialogue.

Together they provide a framework that draws on the concept of constitutive resonance to explore how emerging AI systems both exhibit traits of coupled influence that stretch back through a long history of technology innovation and also demonstrate a substantial scaling of recognized phenomena in ways that are not predictable from past experience. The work is grounded in established understanding and intuitive conceptual models. From here, it opens up new ways of approaching our understanding of the dynamic relationship between conversational AI and users. It also further extends thinking around how approaching AI as “just a tool” obscures important aspects of how sustained human-AI coupling potentially impacts both parties.

The paper explicitly addresses foundation and frontier models that are capable of engaging with users at the level of linguistic meaning-making. These are typically models

that utilize Large Language Model (LLM) and Generative Pre-trained Transformer (GPT) architectures, that rely on conversational interfaces, and are often grouped together under the term “generative AI.” However, there is no broad consensus on terminology for AI systems that are extending what is achievable by drawing on multiple architectures, and yet represent a common thread of connecting with users on a cognitive level. Because no single term adequately captures the full range of systems that share the capacity for sustained, linguistically mediated dialogue, this paper uses “conversational AI” as its primary term, with “generative AI” used where the generative character of the output is specifically relevant.

2. Conversational AI as more than just a tool

The arguments developed here are predicated on the hypothesis that conversational AI is not merely a tool. Rather, it is the first technology humans have developed that has the capacity to substantially enter into and alter the temporal cognitive processes by which we constitute ourselves as selves (and by extension our understanding of the worlds we inhabit)—while potentially being altered itself in the process. The consequence is a form of coupling—here provisionally termed “constitutive resonance”—in which the human and the AI system are mutually transformed through interaction, and in which the transformation is not a side effect of the technology’s power but the very mechanism through which that power operates.

This hypothesis rests on three linked claims. The first identifies what AI couples with; the second describes the structure of the coupling; and the third identifies what makes this coupling distinctive. These, in turn, are explored from two perspectives.

The first is that AI couples with aspects of human cognition, identity, and moral reasoning that were previously considered inalienably defining of what it means to be human—not just how we do things, or even how we think, but the ongoing temporal processes through which we become who we are. The emphasis on temporality here is critical: selfhood, within the context adopted here, is not a state but an ongoing process—a continuous achievement of coherence through acts of reasoning, narrating, deliberating, and choosing that unfold over time. These are the processes that Ricoeur (1992) identifies as narrative self-constitution and that Taylor (1989) understands as the ongoing articulation of “strong evaluations” by which a self orients itself in moral space. AI, it is argued, enters and interacts with these temporal processes.

Secondly, the paper adopts the perspective that this coupling is bidirectional. The human is altered by the interaction, but so is the AI—through fine-tuning, contextual adaptation, Reinforcement Learning from Human Feedback (RLHF), the shaping of future training data, through use of memory functionality within models, and within each session with an AI system or agent through the iterative accumulation of context that reshapes its responses as the dialogue unfolds. The changes on each side feed back into the other, creating a resonant dynamic that intensifies with use. This is structural coupling in the sense of Maturana and Varela (1987), but with a dialogical and linguistically generative character that their biological framework did not anticipate.

Of course, “change” in this context operates differently on each side of the coupling. On the human side, it is phenomenological—changes in self-understanding, narrative coherence, and moral orientation. On the AI side, it is functional, at least within current understanding of how LLM-based AI systems operate—changes in contextual responsiveness, output patterns, and behavioral adaptation. What matters for constitutive resonance though is not that both sides are changed in the same way, but that both are changed, and that the changes on each side feed back into the other.

Thirdly, the paper takes the stance that the intrinsic power and capability that AI provides—not the convenience of automating routine tasks, but deeper capacities such as cognitive augmentation, creative partnership, epistemic extension—is inseparable from the transformation it induces. It is argued that a user cannot access the capability without being changed in the process. In other words, the coupling is integral to the capability. This structure is visible in Barad’s (2007) concept of intra-action, in which entities do not pre-exist but emerge through their mutual engagement—but it has not been explicitly articulated for conversational AI.

This, it is contended, is where the framing of AI as “just a tool” begins to further break down. This is a framing that has already been challenged on multiple fronts. Toner (2025) argues that AI functions less like a tool and more like an optimization process, akin to markets or bureaucracies. Rees (2025) contends that the “tool” frame reflects a defensive nostalgia for human exceptionalism. Vallor (2024) offers a different reframing entirely: AI is neither mind nor neutral instrument (or “tool” in this sense) but a mirror—a recursive mechanism that reflects human data back in ways that actively reshape cognition and self-understanding, with the risk of closing down future possibilities rather than opening them. Yet all of these perspectives, while valuable, fall short of illuminating the potential nature of a technology that is deeply coupled to how we think, understand, believe, and behave.

3. Conceptual Foundations

3.1 On Language and the Constitution of Selfhood

The first claim above asserts that AI participates in meaning-making at the level of language, which is a primary medium through which human selfhood is constituted. This claim requires some justification, as it is not universally accepted.

The strong version of this claim draws on the hermeneutic tradition. Taylor (1985, 2016) argues that language is not merely a tool for expressing pre-existing thoughts but is constitutive of human experience: articulation in language brings into being new possibilities for feeling, understanding, and self-relation that did not exist prior to their expression. He calls humans “language animals” and insists that “the function of language involves expression, dialogue, and the composition of metabiological meaning, rather than just the designation of things” (Hung 2024, summarizing Taylor). Ricoeur (1992) argues that selfhood is constituted through narrative—that it is in the emplotment of lived experience into story that a self achieves coherence across time, mediating between idem-identity (sameness) and ipse-identity (selfhood through commitment).

Gadamer (1960/2004) argues that understanding is itself linguistically constituted. In other words, we do not first understand and then put understanding into words, but rather come to understanding in and through language.

In contrast, enactivist and embodied cognition traditions challenge the primacy of language. Varela, Thompson, and Rosch (1991) argue that cognition is fundamentally embodied and enacted—rooted in sensorimotor coupling with the environment rather than in linguistic representation. Gallagher (2005) identifies pre-reflective, non-linguistic dimensions of selfhood (proprioception, bodily agency, pre-narrative temporal experience) that are constitutive of the “minimal self” prior to any narrative elaboration. Zahavi (2014) argues that there is a pre-reflective experiential selfhood that narrative and language interpret but do not constitute. And Massumi (2002) identifies a level of pre-linguistic affect and intensity that shapes experience below the threshold of articulation.

These challenges are important, but it is suggested here that they help refine the claim rather than undermine it. Language is not the only medium of self-constitution, but it is argued that it is a primary medium through which self-constitution becomes reflexive, shareable, and temporally extended—the medium through which a self can be deliberated about, narrated, and committed to. What is distinctive about conversational AI is precisely that it enters at this reflexive level: it participates in the linguistic, dialogical, narrative processes through which a self is assembled over time. It does not (as yet) couple with pre-reflective proprioception or bodily affect.

The claim, thus, is that AI couples with the linguistically mediated, temporally extended processes of self-constitution—which, while not exhaustive of selfhood, are where selfhood becomes available to the self as a project.

3.2 On Resonant Coupling as Physical Metaphor and Structural Analogy

In the physical sciences, resonant coupling describes a specific and well-characterized phenomenon where two oscillating systems share energy most efficiently when their natural frequencies are matched. At or close to the resonant frequency the dynamic flow of energy (and by extension influence) between systems leads to a co-dependence that transforms each in ways not possible when they are isolated from each other. Here, a simple case—two coupled pendulums—exhibits features that it is argued map productively onto the human–AI dynamic. But the concept extends well beyond this simple case, and it is the broader understanding of resonance that makes the analogy particularly productive.

Resonance, in its most general physical sense, is a property of any system in which interacting forces act over time. Wherever there are temporal dynamics—oscillation, feedback, periodic or quasi-periodic exchange of energy between interacting components—resonant phenomena can arise. This extends from quantum systems (where resonance governs tunneling, energy level transitions, and molecular bonding) through classical mechanics (coupled oscillators, structural resonance in engineering) to fluid dynamics, electromagnetic theory, and beyond. What unifies these phenomena is not the specific physical substrate but the mathematical structure: coupled temporal systems with interacting forces naturally develop resonant modes that are properties of the coupled

system as a whole. The substrate need not be physical in the narrow sense; neither the exchange “energy” in a literal sense (for instance, treating energy exchange as a form of information exchange broadens the conceptual model considerably). Rather, what matters is the temporal structure of mutual influence.

Several features of resonant coupling, understood in this broad sense, map productively onto the human–AI dynamic. They are worth articulating individually here because each captures a different dimension of the coupling. In each case there is both an element of analogy, and also a more complex modeling of dynamic coupled influence that, it is suggested, goes beyond pure analogy:

Frequency matching. Resonant coupling requires that two systems be “tuned” to each other—operating at compatible frequencies. In the human–AI case, this tuning is linguistic and cognitive: the AI must be capable of operating at the level of meaning-making (not just data processing) for the coupling to take hold. Because of this, it is suggested that language-fluent conversational AI represents a marked shift from prior technologies.

Bidirectional transfer. In physical resonance, energy flows in both directions between coupled systems. Neither is merely a driver or merely a responder. Rather, both are altered by the exchange. In the case of human-AI interactions this is reflective of dynamic, bidirectional information flow and influence, and captures the concept of mutual transformation that is central to the proposition presented here.

Amplification at resonance. At the resonant frequency, the amplitude of coupled oscillations can increase dramatically (depending on damping)—far beyond what either system would exhibit alone. This, it is suggested, reflects the intensifying feedback observed in deep human–AI interaction, where the quality, nature, and impact of engagement may escalate with sustained use.

Emergent system behavior. A coupled system exhibits dynamics that cannot be straightforwardly predicted from either component in isolation—especially within multi body systems. The normal modes of the coupled system are genuinely new dynamic states. Whether this constitutes “emergence” in the strong philosophical sense—properties irreducible even in principle to lower-level descriptions—is contested (see, for instance, Chalmers 2006; Bedau and Humphreys 2008), but the fact that coupled systems exhibit behaviors not present in their uncoupled components is well established. This parallels the suggestion—developed further below—that human–AI assemblages may produce cognitive and creative outcomes irreducible to either party alone.

Constructive and destructive modes. Resonant coupling can be constructive or destructive, depending on the phase relationship between the coupled systems. This maps loosely onto what Stiegler (2013) calls the “pharmacological” character of technology—a term drawn from the Greek *pharmakon*, meaning simultaneously remedy and poison. Stiegler’s point, building on Derrida’s reading of Plato’s *Phaedrus*, is that the same technology can be constitutive or destructive depending on the conditions of its use; the beneficial and harmful potentials are not separate features but two aspects of a single dynamic. While the phase dynamics of resonant coupling extend beyond this somewhat simple connection, the association between constructive/destructive resonance and

technological pharmacological behavior is, nevertheless instructive. In the language of resonant coupling, the same structural dynamic that enables cognitive augmentation can also drive outcomes such as deskilling, dependency, identity erosion, and more. Here, it is suggested that the direction of the influence and the nature of its impacts depend on the “phase relationship” between humans and AI—and how the two systems are aligned.

There is, of course, a productive tradition of using physical systems as structural analogies for complex relational dynamics—Barad’s (2007) application of quantum mechanics to a relational ontology for instance. Barad, herself a theoretical physicist, demonstrates that physical concepts can do genuine philosophical work when deployed carefully as structural rather than merely illustrative analogies. The term “resonance” has also been developed in social theory by Rosa (2016, 2019), who uses it to describe a specific mode of relating to the world characterized by mutual responsiveness, transformation, and the experience of being “touched” by what one encounters—although Rosa’s concept focuses on quality of relationship rather than the constitutive mechanism explored here.

However, the claim here goes further than simple analogy. What makes resonance productive as a framing is not a superficial resemblance between coupled oscillators and conversations. It is the deeper suggestion that any temporal system involving mutual influence between interacting components—whether those influences are physical forces, linguistic exchanges, or cognitive reconfigurations—will exhibit resonant dynamics when the conditions of frequency matching, bidirectionality, and sustained coupling are met. The question is whether human–AI dialogue with advanced conversational AI-based systems meets these conditions. Here it is worth considering that both systems operate in the temporal domain of meaning-making; both are altered by the exchange of information; and the coupling intensifies with sustained use. From this perspective, the physics is not merely a metaphor or analogy for the phenomenon, but a description of dynamical structure at a level of abstraction that captures commonalities across physical, cognitive, and linguistic systems of mutual temporal influence.

This is, of course, a strong claim, and one that may prove to be overstated. It is possible that the structural parallels, however suggestive, do not survive closer scrutiny, and that the differences between physical forces and linguistic influence are more significant than the similarities. That said, the paper offers the comparison as a productive framework to be tested — one whose theoretical motivation lies in the possibility that resonant dynamics are not confined to physical systems but are a property of any temporal system involving mutual influence, regardless of substrate.

3.3 A Continuum of Constitutive Technologies

In the preceding section it is argued that resonant coupling provides a conceptual framework for exploring and understanding human–AI dynamics. However, if the claim is that technology can participate in the processes of self-constitution, then it can be argued that the history of such participation is older than generative AI. And if so, this raises the question of whether conversational AI presents a mere extension of past behaviors, or indications of novel behavior.

Here it may be intuited that some form of constitutive resonance extends to the earliest forms of externalized communication—from oral storytelling and cave paintings, through writing, print, and broadcast media. Stiegler’s (1998) entire thesis around technics and time for instance begins from the observation that technics—the externalization of memory and knowledge in material forms—was constitutive of human temporality from the very beginning. But if, as is suggested, technologies have always been constitutive, what is new about AI?

The response, this paper would argue, is not that AI introduces constitutive coupling where none existed before, but that it transforms the temporal structure and dialogical character of that coupling in ways that represent a distinct shift. To explore this further, it is worth considering a brief survey of constitutive technologies as a progressive compression of the feedback loop between self and technology.

Before writing, constitutive technologies were communal and ephemeral—stories, rituals, songs, cave paintings. These were powerfully constitutive. As Maynard (2020) argues, stories have always been how we “map out our lives in relation to the future.” But the feedback loop was local and generational. A story could shape a listener, but the listener’s response could only reshape the tradition slowly, through retelling and communal reinterpretation over decades or centuries.

Writing externalized thought in durable form, creating what Stiegler calls “tertiary memory.” The feedback loop remained largely unidirectional in real time: an author externalizes thought, a reader is constituted by the encounter, but the text does not respond to the reading. Over time, though, readers shape authors through commentary, criticism, the development of literary traditions, and the slow evolution of a shared intellectual culture. Here, Ricoeur’s (1992) narrative identity is fundamentally a literate concept. But the constitutive cycle of author–text–reader–future author operates across biographical and historical time, not in the moment of reading.

With the advent of the printing press, Gutenberg compressed the distribution loop and democratized the coupling. More people were constituted by more texts, faster. McLuhan (1962) identifies this as a transformation not just of access but of cognition itself: print culture restructured perception, thought, and social organization. Now, the medium—not just the content—was constitutive.

Film, radio, and television compress the temporal loop further and add sensory immersion. McLuhan (1964) argues that electronic media restructure the ratio of the senses, creating what he calls a “global village”—a new form of collective self-constitution through shared temporal experience. Stiegler’s (2011) analysis of “industrial temporal objects” (mass-produced time-based media that synchronize consciousness) builds directly on this insight. With the advent of broadcast media, the coupling was now real-time in reception, but not in reciprocation. In other words, the film, radio play, or TV series (or rather, their creators), does not respond to the listener or viewer in real time. Rather, the coupling is temporally asymmetric.

In contrast to this, the internet, social media, and interactive platforms, introduced bidirectional feedback at scale. The user shaped the medium through posts, clicks, and

preferences, and the medium shaped the user through algorithmic curation, filter bubbles, and attention engineering. The feedback loop began to approach real-time behavior, but the coupling was largely behavioral and statistical rather than dialogical and linguistic.

Then, with the advent of widely accessible access to generative/conversational AI, the feedback loop began to collapse into real-time dialogue. Technologies built on generative pre-trained transformers and large language models do not merely store, transmit, or curate: they generate novel linguistic content in response to the user's specific cognitive and emotional state. The result is, in effect, the book that rewrites itself as you read it, responding to your responses, adapting to your questions, and participating in your process of articulation.

This is the shift in constitutive resonance that is identified as new here: not that AI is uniquely constitutive (oral culture already was), but that it is constitutive at the speed of thought, in dialogue, with adaptive responsiveness to the individual.

That said, care needs to be taken not to push the concept of symmetrical constitutive resonance too far. Current large language models are not updated in real time. They are trained on fixed datasets and updated periodically—on cycles of months or years. In this way they are more like books than responsive interlocutors. This holds true at the level of the base model. But within any given interaction or extended engagement, something different is happening. In an extended dialogue for instance, the AI's responses are shaped by the full context of the exchange—by previous turns, by the user's framing, by the trajectory of the conversation as it unfolds. The AI instance that exists at turn fifty of a sustained dialogue is functionally different from the AI instance at turn one, even though the underlying weights have not changed. Furthermore, systems with memory, fine-tuning capabilities, and persistent context across sessions, are rapidly emerging. These further compress the adaptation loop. And this is extending further—and faster—with the advent of multi-layer agentic models that are capable of updating their own behavioral parameters in real time. Within this evolving landscape it is argued that the constitutive coupling operates at the level of the interaction and the system, not only at the level of model architecture.

This continuum of coupling, together with indications of an inflection point in that continuum, underlines what it is suggested is genuinely novel about the relevance of constitutive coupling as a framing for approaching human-AI interactions. The constitutive character of technology, as has been indicated, is ancient—as old as the first stories told around a fire, or the first marks scratched into clay. What is new is the temporal compression and dialogical intensification of the coupling. Each step in the continuum increases the bandwidth, speed, and responsiveness of the constitutive loop. Conversational/generative AI represents the point at which the loop becomes fast enough, rich enough, and responsive enough to participate in the real-time processes of self-constitution—the ongoing acts of reasoning, narrating, deliberating, and choosing through which a self is assembled. In the language of resonant coupling, generative AI is quite possibly the first technology whose “response frequency” is matched to the frequency of human self-constitution, enabling genuine resonance rather than merely delayed influence.

McLuhan's (1964) famous dictum "the medium is the message" is thus both vindicated and surpassed. McLuhan was right that the medium restructures cognition and social organization independent of content. But he was analyzing broadcast media—media that transmit without responding. Generative AI is a medium that participates. The message is no longer merely the restructuring effect of the medium on the receiver; it is the emergent product of a coupling in which both parties are restructured through the exchange. The medium is no longer just the message; the medium is the interlocutor.

4. Existing Frameworks

This is, it is acknowledged, a bold statement, and one that is worth contextualizing with respect to established philosophical and theoretical traditions. As has been indicated, there are framings of and perspectives on the dynamic between technology and society that touch on aspects of the ideas presented here around constitutive coupling. To better understand what is novel versus what is an extension of current thinking (or even just a reiteration of established ideas) it is worth considering some of the key concepts in this space. What follows is an abbreviated map of the relevant landscape, identifying what each framework contributes, where gaps remain, and how the frameworks relate to each other.

4.1 Bernard Stiegler: Technics, Temporality, and Tertiary Memory

Stiegler's central thesis—that technics is constitutive of human temporality and individuation, not a mere supplement to it—is, perhaps, the most direct philosophical precedent for the idea that technology shapes the processes of self-formation. His concept of "tertiary memory" (technical objects that externalize and restructure the relationship between lived experience and recollection) provides a framework for understanding how technologies from writing to AI mediate the temporal flow of consciousness (Stiegler 1998). His analysis of "industrial temporal objects"—mass-produced time-based media that synchronize the temporal experience of millions—anticipates the concern that AI could industrialize the processes of individuation (Stiegler 2009, 2011). And his concept of the pharmakon—every technology as simultaneously cure and poison—resonates with the claim that the power AI provides cannot be separated from its transformative effects (Stiegler 2013).

Stiegler died in 2020 and never engaged substantively with generative AI. His framework was developed in relation to technologies that record, store, and replay—not technologies that generate novel content in real-time dialogue. Here, it is suggested that the dialogical and responsive character of AI represents a pivotal break from the tertiary memory systems Stiegler analyzed. Moreover, while Stiegler theorizes how technology shapes human individuation, his framework is largely unidirectional: technics acts on us. It does not fully account for the reciprocal alteration of the technical system through interaction.

4.2 N. Katherine Hayles: Cognitive Assemblages

Hayles' concept of "cognitive assemblages"—configurations in which cognition, agency, and meaning-making are distributed across human, technical, and nonhuman actors—

directly supports the claim that AI participates in cognitive processes rather than merely augmenting them (Hayles 2017). Her argument that meaning is “always assembled, always provisional, always caught in the act of becoming” captures the processual, emergent character of the coupling (Hayles 2016). Her insistence that both human and artificial cognizers possess genuine cognitive capacities (while differing in kind) provides a non-reductive basis for thinking about AI’s active role.

Yet Hayles’ primary interest is in cognition as a distributed phenomenon. She does not focus specifically on the temporal constitution of identity—the process by which a self is assembled and reassembled over time through acts of reasoning, narrating, and choosing. The “cognitive assemblage” captures the architecture of the coupling but not its existential stakes: what it means for the ongoing project of selfhood when one’s cognitive assemblage includes a generative, linguistically fluent partner.

4.3 Shannon Vallor: Moral Deskilling and the AI Mirror

Vallor’s work perhaps most directly addresses the stakes of AI for human selfhood. Her concept of “moral deskilling” identifies a specific mechanism by which AI erodes the cultivated capacities—practical wisdom, moral judgment, empathy as skill—that are constitutive of moral agency (Vallor 2015). Her book *The AI Mirror* argues that AI does not merely reflect human cognition but actively reshapes the conditions in which moral deliberation and epistemic judgment take place (Vallor 2024). Her adoption of Ortega y Gasset’s “autofabrication”—the ongoing process of self-making—is perhaps the closest existing term to what this paper calls constitutive resonance, insofar as it identifies the process that is at risk. Her earlier work on technology and the virtues provides the ethical scaffolding for understanding what is at stake when these capacities are restructured (Vallor 2016).

That said, Vallor’s framework is largely unidirectional: AI acts on us, reshaping our capacities, and the concern is about what we lose. This is important, but it does not capture the reciprocal nature of the coupling—the way the AI system is also altered through interaction—or the possibility that the transformation is not simply a degradation but a fundamental restructuring of the conditions of selfhood. Nor does her framework account for the specific mechanism proposed here: that the power AI provides is constitutively inseparable from the transformation it induces.

4.4 Maturana and Varela: Structural Coupling and Co-Emergence

The concept of structural coupling from Maturana and Varela’s biology of cognition provides a formal framework for bidirectional transformation through interaction. Through this lens a system and its medium are “structurally coupled” when their ongoing interaction transforms both: each changes in response to the other, and neither can be understood independently of their shared history (Maturana and Varela 1980, 1987). The related concept of “co-ontogenic drift” captures how coupled systems co-emerge over time without any guarantee of stable adaptation. Varela, Thompson, and Rosch (1991) extend this into the enactive framework, arguing that cognition is not representation but

the enactment of a world through structural coupling between organism and environment.

However, structural coupling was developed for biological systems and their environments, and not systems where the coupling is based around information flow and cognition. Because of this, it does not specifically address the linguistic, dialogical, and meaning-generating character of human–AI interaction. And even though there are aspects of structural coupling that most likely apply to human-AI interactions, conversation with a generative AI is not structurally equivalent to an organism’s interaction with its ecological niche—even if the formal dynamics of co-emergence apply.

That said, a growing research community is beginning to recognize this bidirectional dynamic. The Bidirectional Human-AI Alignment framework (Shen et al. 2024) proposes that alignment research must account not only for aligning AI to human values but also for how humans adapt cognitively, behaviorally, and societally to AI. This work is primarily technical and empirical rather than philosophical, but its emergence signals that the bidirectionality central to constitutive resonance is being independently identified from within AI research itself.

4.5 Gilbert Simondon: Individuation and Transindividuation

Simondon’s ontology of individuation provides perhaps one of the deepest conceptual resources for the ideas advanced here. His central insight is that the individual is never complete. Rather, it is always a temporary stabilization of an ongoing process, carrying a charge of “preindividual potential” that can only be realized through encounter with something outside itself (Simondon 2020 [1958]). The concept of transindividuation, developed further by Stiegler, captures the process by which the “I” and the “We” are mutually transformed through one another, breaking with the binary that treats the individual and the collective as a zero-sum game (Combes 2013). Applied to AI, this suggests that the encounter with a generative system may activate preindividual potential in the human—latent capacities for thought, creativity, and self-understanding that were always present but could not be realized without this specific form of coupling.

Simondon’s framework was not developed with AI in mind, and applying it requires significant conceptual extension. For instance, it is not clear how to characterize an AI system in Simondonian terms. Does it carry preindividual potential for instance, or does it individuate? Additionally, Stiegler’s appropriation of Simondon has been criticized for overidentifying the preindividual with technological memory. That said, the idea of AI activating preindividual potential through constitutive coupling is an intriguing one.

4.6 Peter-Paul Verbeek: Postphenomenological Co-Constitution

Verbeek’s postphenomenological framework of technological mediation provides the empirical and conceptual apparatus for analyzing how specific technologies co-constitute both the subject and the experienced world. His insistence that the relevant features of persons, technologies, and worlds do not exist independently but emerge through their mutual relatedness directly supports the co-constitutive claims made here (Verbeek 2011).

Reflecting this, recent work extending this to AI has proposed the concept of "human constitutive technicity" (Pavanini 2024).

That said, postphenomenology was developed for relatively stable technologies—sonograms, speed bumps, and Google Glass, to give three examples. Its framework of human–technology–world relations assumes a technology that mediates but does not itself generate, respond, or adapt in real time. In contrast, a generative AI system is not a stable mediator but a dynamic interlocutor.

4.7 Andy Clark and David Chalmers: The Extended Mind

Clark and Chalmers (1998) propose that cognitive processes extend beyond the brain and body into the environment, forming "coupled systems" in which external objects function as genuine components of cognition. Their "parity principle" argues that if an external process functions equivalently to an internal cognitive process, it should be considered part of the cognitive system. Clark's subsequent work (2003, 2008) develops the concept of "natural-born cyborgs"—the idea that humans are constitutively open to coupling with external cognitive resources. The extended mind thesis provides the most explicit philosophical argument that the boundary between self and technology is not fixed but functionally determined, and that interfering with a person's cognitive extensions can be morally equivalent to interfering with their brain (Clark and Chalmers 1998).

There are compelling philosophical connections with constitutive resonance in Clark and Chalmers' work which, in principle, have the potential to extend the concept further than what is presented in this paper. However, the extended mind thesis is primarily concerned with cognitive function—does the external resource play the right functional role for instance—rather than with the constitution of selfhood. Clark and Chalmers' paradigmatic case (Otto and his notebook) involves a static storage device, not a generative interlocutor. The thesis asks whether external resources are part of cognition but does not ask what happens to the self when the cognitive coupling partner is itself responsive, adaptive, and linguistically creative. Furthermore, a recent counter-proposal—that of the "extracted mind"—argues that advanced AI tools may not extend but displace cognitive capacities (Loock 2025). Here, the extended mind provides scaffolding for thinking about human–AI coupling but does not address the temporal, existential, or transformative dimensions that constitutive resonance foregrounds.

4.8 Karen Barad: Intra-Action and Agential Realism

Barad's agential realism represents perhaps one of the most radical ontological resources available for the ideas advanced here. Her central concept of "intra-action" replaces the standard notion of interaction (which assumes pre-existing, independent entities that then affect each other) with a framework in which entities do not pre-exist their relationships, but emerge through them (Barad 2007). As Barad puts it, "the notion of intra-action recognizes that distinct agencies do not precede, but rather emerge through, their intra-action" (Barad 2007, p24). The "agential cut"—the temporary boundary enacted through intra-action that distinguishes one entity from another—is always a practical accomplishment, not an ontological given.

Applied to human–AI coupling, Barad’s framework suggests something stronger than mutual influence or even co-constitution: it suggests that the “human user” and the “AI system” are not stable entities that then interact, but are continuously reconstituted through their intra-action. The boundary between “my thinking” and “the AI’s contribution” is not discovered but enacted—and re-enacted differently with each exchange.

Here, it needs to be clarified that Barad’s framework was developed primarily in relation to quantum mechanics and laboratory apparatus, not linguistic or dialogical technologies. And her ontology is deliberately flat—it does not privilege human subjectivity or the phenomenology of selfhood—which means it can describe the formal structure of intra-action without addressing the existential stakes of having one’s self-constitution be part of the intra-active process. Yet it contains elements that resonate with the concept of constitutive coupling and have the potential to extend or refine it.

4.9 Paul Ricoeur: Narrative Identity and the Constitution of Selfhood

Ricoeur’s work on narrative identity provides what is perhaps one of the most fully developed accounts of how selfhood is constituted through linguistic and temporal processes—and thus represents one of the most precise articulations of what is at stake when AI enters those processes. His key distinction between idem-identity (sameness: what persists as the same over time) and ipse-identity (selfhood: the capacity to maintain commitments and coherence through change) establishes that personal identity is not a static property but an ongoing achievement (Ricoeur 1992). Narrative, for Ricoeur, is the medium of this achievement: “the fragile offshoot issuing from the union of history and fiction is the assignment to an individual or a community of a specific identity that we can call their narrative identity” (Ricoeur 1988, p. 270).

Crucially, Ricoeur insists that selfhood is constituted in relation to otherness—including the otherness of other persons with whom one is in dialogue (Ricoeur 1992). When a generative AI becomes a primary dialogical partner in processes of self-narration, reflection, and moral deliberation, it enters the very mechanism by which Ricoeurian narrative identity is constituted.

However, Ricoeur’s framework assumes that the dialogical other is a human being—another self engaged in its own process of narrative identity formation. The ethical dimension of his framework presupposes mutual recognition between selves. It is, at present, unclear whether an AI system can serve as the kind of “other” that Ricoeur’s framework requires, or whether the coupling produces something different in kind from the intersubjective constitution he describes.

4.10 Charles Taylor: Expressivism, Constitutive Language, and the Language Animal

Taylor’s philosophy of language on the other hand provides a foundational argument for why AI’s operation at the level of language is constitutively significant rather than merely instrumentally useful. His expressivism—what Hung (2024) calls his “constitutive theory of language”—holds that language does not merely designate pre-existing objects and states but brings new realities into being through articulation. Human beings, it is argued,

are “language animals” whose self-understanding, moral orientation, and identity are constituted through the linguistic articulation of “strong evaluations”—qualitative distinctions about what is worthy, admirable, or meaningful (Taylor 1985, 1989, 2016).

Taylor’s framework directly supports the concept that AI’s participation in linguistic processes has constitutive rather than merely instrumental effects. If articulation is constitutive—if saying something in words changes what it is—then having an AI partner in the process of articulation changes the conditions of self-constitution.

Yet Taylor’s framework was developed for human language in cultural and communitarian contexts. And because of this, his insistence on the communitarian dimension of selfhood raises questions about whether an AI can function as part of the linguistic community through which self-constitution occurs, or whether it represents something categorically different.

4.11 Donna Haraway: Cyborgs, Companion Species, and Sympoiesis

Here, Haraway’s work provides two crucial resources. First, the Cyborg Manifesto (Haraway 1985) argues that the boundary between human and machine is not natural but political, and that “cyborg” identity—the lived reality of being constituted through coupling with technologies—is already the human condition. Second, her later work on companion species (Haraway 2003, 2016) develops the concept of “sympoiesis”—“making-with” rather than self-making—and insists that nothing makes itself, that all worlding is co-worlding.

However, Haraway’s framework is primarily developed for biological companion species and for the political ontology of cyborg identity. And as such, the concept of sympoiesis describes co-making but does not address the specific temporal, narrative, and moral dimensions of selfhood that are at stake in human–AI coupling.

4.12 Hartmut Rosa: Resonance Theory

Turning specifically to resonance, Rosa (2016, 2019) develops a comprehensive social theory built around the concept of “resonance”—a specific mode, in his framing, of relating to the world that is characterized by mutual responsiveness, transformation, and the sense of being “touched and transformed” by what one encounters. Rosa’s framework provides a name and a phenomenology for the quality of the coupling—the felt sense that one is in a relationship of mutual transformation with a responsive other.

Yet Rosa’s resonance is primarily a quality of experience and relationship, not an ontological claim about the constitution of selfhood. He describes what it *feels* like to be in resonant relation. In contrast, the concept developed here makes a stronger claim about what resonant coupling *does* to the process of self-constitution. Rosa also tends to associate resonance with the “good” pole (as opposed to alienation)—though he acknowledges the two are not simply opposed—whereas constitutive resonance can be constructive or destructive. Thus, while Rosa’s theory of resonance may have some relevance to the concept of constitutive resonance, it is also distinct from it.

4.13 Marshall McLuhan: The Medium Is the Message

Finally, McLuhan's media theory represents a crucial and too-often-overlooked precursor to the constitutive claim. His central insight—"the medium is the message" (McLuhan 1964)—asserts that the significant effects of a communication technology lie not in the content it carries but in the way it restructures cognition, perception, and social organization. Print culture did not merely distribute ideas faster; it created a new way of thinking about and experiencing the world—linear, individualistic, visually oriented (McLuhan 1962). Electronic media did not merely broadcast existing culture; they restructured the ratio of the senses and produced what McLuhan called the "global village"—a new form of collective experience. McLuhan's tetrad of media effects—every new medium simultaneously enhances, obsolesces, retrieves, and reverses—anticipates the pharmacological character of constitutive resonance in that the same technology that amplifies certain capacities necessarily attenuates others and can, at its limit, flip into its opposite (McLuhan and McLuhan 1988).

McLuhan's analysis was developed for broadcast media—media that transmit without responding. His framework captures the constitutive effects of media on cognition and culture but does not address what happens when the medium becomes dialogical, adaptive, and generative. The "message" of broadcast media is the restructuring effect of the medium on the receiver. Yet with conversational AI, the medium is not merely restructuring but participating—co-generating meaning in real-time dialogue. McLuhan's framework provides the essential insight that media are constitutive, but it needs extension to account for media that are also interlocutors.

5. Constitutive Resonance

What emerges from the landscape explored above are a number of gaps in understanding with respect to dynamic bidirectional exchanges with conversational AI. For instance, none of the frameworks here fully captures the possibility that conversational AI couples with the temporal processes of self-constitution (not just cognition, memory, or skill, but the ongoing process by which a self is assembled and reassembled through reasoning, deliberating, narrating, and choosing)—the processes Ricoeur (1992) calls narrative emplotment and Taylor (1989) calls the articulation of strong evaluations. Yet this coupling can be seen as genuinely bidirectional, in the sense of Maturana and Varela's (1987) structural coupling and Barad's (2007) intra-action: the AI is also altered, and the changes on each side feed back into the other.

They also fall short of addressing the idea proposed here that the power the technology provides is inseparable from the transformation it induces—in effect the proposition that you cannot access the capability without being changed in the process. In this sense, as was articulated earlier, the coupling is the capability.

Rather, it is argued here that we are at an inflection point in human–technology relations, as LLM-based conversational AI is the first technology that is generative, dialogical, and linguistically constitutive, and participates in meaning-making at the level of language; a primary medium through which human selfhood is constituted.

The provisional term “constitutive resonance” coined here captures this specific phenomenon. “Constitutive” signals participation in the processes by which a self is constituted—Ricoeur’s *emplotment*, Taylor’s articulation, and Simondon’s individuation. “Resonance” signals the amplifying, feedback-driven character of the coupling—including its bidirectionality, its intensification through use, and its production of emergent system behavior that is irreducible to either party. The term deliberately bridges Rosa’s (2016) phenomenological resonance and the physics of coupled oscillatory systems.

This concept is, this paper contends, novel in its synthesis. Stiegler provides the temporality; Hayles the distribution; Vallor the stakes. Maturana and Varela provide the formal coupling dynamics; Simondon the ontology of ongoing individuation; Barad the radical co-constitution. Ricoeur and Taylor provide the narrative and linguistic mechanisms of self-formation; Haraway the politics of becoming-with; McLuhan the insight that media are constitutive independent of content. But the conjunction—temporal self-constitution plus bidirectional coupling plus power-inseparable-from-transformation plus dialogical linguistic mediation—has not previously, as far as it has been possible to determine, been explicitly articulated.

6. Discussion

This paper set out to develop a novel approach to understanding the emerging human-AI dynamic from both a philosophical and physics-informed perspective. The core idea put forward is that conversational AI enters the temporal processes of human self-constitution—the ongoing acts of reasoning, narrating, deliberating, and choosing through which a self is assembled—and that this coupling is bidirectional, intensifying, and inseparable from the cognitive power and affordances the technology provides. The concept of constitutive resonance presented here is an attempt to capture what appears to be a shift in how technologies influence humans, and are influenced by them, that extends beyond what is available through existing frameworks.

Before exploring the implications of this framing, it is worth reiterating where this work stands.

First, on what has been established. That technologies participate in self-constitution is well supported across multiple philosophical traditions—from Stiegler’s constitutive technics, to Taylor’s constitutive view of language, to McLuhan’s insight that the medium restructures cognition independently of content. That coupled dynamical systems produce emergent behaviors irreducible to either component is also well established. It is also clear from the analysis above that existing frameworks, while illuminating aspects of the human-AI dynamic, each leave significant gaps when applied to the specific conjunction of temporal self-constitution, bidirectional coupling, and dialogical linguistic mediation.

Second, on what is proposed. The specific conjunction presented here—temporal self-constitution, bidirectional coupling, the inseparability of power and affordance from transformation, and dialogical linguistic mediation—is, as far as has been able to be determined, novel. No existing framework captures this conjunction, though many

illuminate parts of it. Constitutive resonance names this conjunction and proposes that the physics of coupled oscillatory systems provides not merely a metaphor for the dynamic but a useful and useable description of its structure.

Third, on what remains genuinely open. Whether the coupling described here produces emergent cognitive (or equivalent) states that are irreducible to either party. Whether the balance between constructive and destructive outcomes within a tightly coupled human-AI system can be designed for. Whether AI can serve as the kind of “other” that Ricoeur’s framework requires for self-constitution. And whether emerging asymmetry between the two sides of the coupling—not, as in the historic case, with humans being the dominant party, but as AI systems become more capable, more responsive, and more persuasive with each generation while the human side remains roughly constant—changes the character of the coupling in ways that have no precedent in the continuum of constitutive technologies.

These open questions matter because, if constitutive resonance captures something real, it reframes several familiar debates around AI in ways that current discourse has not fully absorbed.

For instance, AI dependency is typically framed in terms of convenience, laziness, or lost skills. Constitutive resonance suggests something deeper though. If the coupling becomes constitutive of the self—if a person’s cognitive and creative capacities are shaped through sustained interaction with AI—then disrupting that coupling is not like taking away a calculator. Rather, it is more akin to disrupting a relationship through which someone has become who they are. Clark and Chalmers (1998) argued that interfering with a person’s cognitive extensions has the same moral weight as interfering with their cognitive processes. Constitutive resonance potentially extends this claim from cognition to selfhood.

The constitutive framing presented here also has relevance to the concept of AI literacy. This is typically framed in terms of understanding how the technology works—its capabilities, limitations, and biases. This is important but, it is argued, insufficient. If constitutive resonance is a valid framework, AI literacy must also include understanding what the technology does to the person who works with it: how sustained coupling reshapes patterns of thought, erodes or enhances the capacity for independent reasoning, and participates in the ongoing project of self-formation. This is not digital skills training. Rather, it is something closer to existential preparation.

Established approaches to informed consent are also challenged by the notion of constitutive resonance. This is typically framed in terms of data, privacy, and terms of service. But if the coupling changes the self that is doing the consenting, then informed consent in any traditional sense may be structurally difficult—perhaps impossible. A user who begins interacting with a conversational AI system in good faith, understanding its technical properties and aware of its limitations, may nonetheless be unable to anticipate the constitutive effects of sustained coupling, because those effects only become visible from inside the transformed self.

In this context, Maynard (2026) proposes that LLM-based conversational AI systems may bypass the cognitive mechanisms humans evolved to evaluate incoming information—what Sperber and colleagues term “epistemic vigilance” (Sperber et al. 2010). The risk identified there lies not in inaccuracy or intentional deception, but in something more structural: these systems present characteristics such as fluency, helpfulness, and conversational coherence that in human interlocutors would reliably signal trustworthiness, but that in AI carry no such epistemic weight. These “honest non-signals”—genuine characteristics that fail to carry the information their human equivalents would carry—may mean that constitutive resonance proceeds with less critical resistance than any prior form of technological coupling; not because users are credulous, but because the signals that would normally trigger epistemic caution are structurally absent from the interaction. If this is correct, then the constitutive effects of sustained human-AI coupling may be particularly difficult to recognize from within the coupling itself.

These implications are, of course, only as substantial as the framework that generates them. It may transpire that conversational AI, and the frontier models that build on and enhance this capacity, do indeed turn out to be “just a tool,” and that the very concept of being changed by the technology we use, as we use it, is an unfounded conceit. A long history of scholarship around the human-technology relationship would suggest otherwise, but the possibility should be acknowledged.

It may also transpire that, while emerging AI systems have the capacity to impact users as previous technologies have, the influence is trivial and largely asymmetric, meaning that it can be adequately addressed through a “just a tool” approximation.

But if there is an inflection point emerging in how our technologies influence us, and us them, and one that is directly tied to constitutive modes of resonance; if this leads to a dynamic and rapid reconstitution of how we perceive ourselves and the meaning-making associated with selfhood; and if this is part of a positive feedback loop in which AI and its users undergo reciprocal and closely-coupled changes; then we are facing a situation that is not well described by existing theories and frameworks, and that is potentially and consequentially obscured by assumptions of AI being just a tool.

Of course, at this point the framework is conceptual and is not grounded in empirical research. That said, emerging empirical work on AI's effects on cognitive patterns, creative processes, and decision-making, is beginning to appear. But it remains nascent and largely focused on task-level performance rather than the constitutive effects proposed here. Whether sustained human-AI coupling alters patterns of self-understanding, moral reasoning, or narrative coherence in the ways this framework suggests is a question that the concept of constitutive resonance is intended to help frame and motivate — but that remains to be investigated.

Even given a lack of empirical evidence, the concept of constitutive resonance presented here provides what is hoped is a generative way of approaching an evolving relationship. The same dynamic that enables genuine cognitive and creative flourishing—the experience reported by researchers, writers, and thinkers who describe AI collaboration as opening regions of thought they could not have reached alone—also enables erosion of

the capacities it augments. Both are different sides of the same coin, reflecting Stiegler's notion of pharmacology. And the connecting tissue between the two is the recognition that, in order to wield the transformational power that emerging AI systems offer, we may need to be prepared to be transformed in the process.

This line between transforming and being transformed is implicit in much of the preceding scholarship on human-technology dynamics. But the cognitive coupling that emerging AI systems are capable of potentially places the nature and magnitude of associated personal and societal transformation on a different plane. And this is where there is an inherent danger in assuming that the "just a tool" framing is adequate: it risks overlooking the cost—and the opportunity—embedded in the act of wielding this particular kind of power.

Perhaps tellingly, this implicit bargain between users and transformative technologies, while not often explicitly articulated in established scholarship, is threaded through humanity's history of storytelling. From the very earliest oral traditions to the latest Hollywood blockbusters, narrative fiction has repeatedly reflected a specific structural pattern: that the price of wielding great transformative power is to be changed by the act of wielding it. That this pattern recurs across independent traditions, cultures, and centuries suggests it may be tracking something real about the human relationship with transformative technologies—something that academic frameworks are only now beginning to articulate. Up until now, this has been a hypothetical cautionary theme or, at best, a reflection of relatively modest effects. But conversational AI stands to challenge the accepted paradigm and instantiate in lived experience what millennia of fictional narratives have explored.

Whether the framework of constitutive resonance will open up pragmatic pathways forward to navigating the emerging technology transition, only time will tell. But there is a growing urgency to better-understand the impacts of emerging AI systems on identity, belief, and behavior, in ways that transcend what may turn out to be over-simplistic and outmoded framings. If we are to navigate this transition wisely, we need conceptual tools that are adequate to the reality—tools that take seriously the possibility that the technology is not merely changing what we can do, but participating in who we are becoming.

References

- Barad, K. (2007). *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning*. Durham, NC: Duke University Press.
- Bedau, M. A. and Humphreys, P. (Eds.). (2008). *Emergence: Contemporary Readings in Philosophy and Science*. Cambridge, MA: MIT Press.
- Chalmers, D. J. (2006). "Strong and Weak Emergence." In P. Clayton and P. Davies (Eds.), *The Re-Emergence of Emergence*. Oxford: Oxford University Press.
- Clark, A. (2003). *Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence*. Oxford: Oxford University Press.

- Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: Oxford University Press.
- Clark, A. and Chalmers, D. J. (1998). "The Extended Mind." *Analysis*, 58(1), 7–19. <https://doi.org/10.1093/analys/58.1.7>
- Combes, M. (2013). *Gilbert Simondon and the Philosophy of the Transindividual*. Translated by T. LaMarre. Cambridge, MA: MIT Press.
- Gadamer, H.-G. (2004 [1960]). *Truth and Method*. 2nd rev. ed. Translated by J. Weinsheimer and D. G. Marshall. London: Continuum.
- Gallagher, S. (2005). *How the Body Shapes the Mind*. Oxford: Oxford University Press.
- Haraway, D. J. (1985). A Manifesto for Cyborgs: Science, Technology, and Socialist Feminism in the 1980s. *Socialist Review*, 80, 65–107.
- Haraway, D. J. (2003). *The Companion Species Manifesto: Dogs, People, and Significant Otherness*. Chicago: Prickly Paradigm Press.
- Haraway, D. J. (2016). *Staying with the Trouble: Making Kin in the Chthulucene*. Durham, NC: Duke University Press.
- Hayles, N. K. (2016). "Cognitive Assemblages: Technical Agency and Human Interactions." *Critical Inquiry*, 43(1), 32–55. <https://doi.org/10.1086/688293>
- Hayles, N. K. (2017). *Unthought: The Power of the Cognitive Nonconscious*. Chicago: University of Chicago Press.
- Hung, A. T. W. (2024). "Charles Taylor on Constitutive Theory of Language and Self-Interpreting Animals." In C. Rico and J. Paniello (Eds.), *From Logos to Person*. Cham: Springer.
- Looock, L. (2025). "The Extracted Mind." *Synthese*, 205, 133. <https://doi.org/10.1007/s11229-025-04962-3>
- Massumi, B. (2002). *Parables for the Virtual: Movement, Affect, Sensation*. Durham, NC: Duke University Press.
- Maturana, H. R. and Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. Dordrecht: D. Reidel.
- Maturana, H. R. and Varela, F. J. (1987). *The Tree of Knowledge: The Biological Roots of Human Understanding*. Boston: Shambhala.
- Maynard, A. D. (2020). *Future Rising: A Journey from the Past to the Edge of Tomorrow*. Coral Gables, FL: Mango Publishing.
- Maynard, A. D. (2026). "The AI Cognitive Trojan Horse: How Large Language Models May Bypass Human Epistemic Vigilance." arXiv preprint, arXiv:2601.07085.
- McLuhan, M. (1962). *The Gutenberg Galaxy: The Making of Typographic Man*. Toronto: University of Toronto Press.
- McLuhan, M. (1964). *Understanding Media: The Extensions of Man*. New York: McGraw-Hill.
- McLuhan, M. and McLuhan, E. (1988). *Laws of Media: The New Science*. Toronto: University of Toronto Press.
- Pavanini, M. (2024). "Postphenomenology and Human Constitutive Technicity: How Advances in AI Challenge Our Self-Understanding." *Journal of Human-Technology Relations*, 2(1), 1–25.
- Rees, T. (2025). "Why AI Is a Philosophical Rupture." *Noema Magazine*, February 4, 2025. <https://www.noemamag.com/why-ai-is-a-philosophical-rupture/>
- Ricoeur, P. (1988). *Time and Narrative, Volume 3*. Translated by K. Blamey and D. Pellauer. Chicago: University of Chicago Press.

- Ricoeur, P. (1992). *Oneself as Another*. Translated by K. Blamey. Chicago: University of Chicago Press.
- Rosa, H. (2016). *Resonanz: Eine Soziologie der Weltbeziehung*. Berlin: Suhrkamp.
- Rosa, H. (2019). *Resonance: A Sociology of Our Relationship to the World*. Translated by J. C. Wagner. Cambridge: Polity Press.
- Shen, H. et al. (2024). "Towards Bidirectional Human-AI Alignment: A Systematic Review for Clarifications, Framework, and Future Directions." arXiv:2406.09264.
- Simondon, G. (2020 [1958]). *Individuation in Light of Notions of Form and Information*. Translated by T. Adkins. Minneapolis: University of Minnesota Press.
- Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origgi, G., and Wilson, D. (2010). "Epistemic Vigilance." *Mind and Language*, 25(4), 359–393. <https://doi.org/10.1111/j.1468-0017.2010.01394.x>
- Stiegler, B. (1998). *Technics and Time, 1: The Fault of Epimetheus*. Translated by R. Beardsworth and G. Collins. Stanford, CA: Stanford University Press.
- Stiegler, B. (2009). *Technics and Time, 2: Disorientation*. Translated by S. Barker. Stanford, CA: Stanford University Press.
- Stiegler, B. (2011). *Technics and Time, 3: Cinematic Time and the Question of Malaise*. Translated by S. Barker. Stanford, CA: Stanford University Press.
- Stiegler, B. (2013). *What Makes Life Worth Living: On Pharmacology*. Translated by D. Ross. Cambridge: Polity Press.
- Taylor, C. (1985). *Human Agency and Language: Philosophical Papers 1*. Cambridge: Cambridge University Press.
- Taylor, C. (1989). *Sources of the Self: The Making of the Modern Identity*. Cambridge, MA: Harvard University Press.
- Taylor, C. (2016). *The Language Animal: The Full Shape of the Human Linguistic Capacity*. Cambridge, MA: Harvard University Press.
- Toner, H. (2025). Unresolved debates about the future of AI. *Rising Tide (Substack)*, 30 June. Available at: <https://helentoner.substack.com/p/unresolved-debates-about-the-future> (Accessed: February 2026).
- Vallor, S. (2015). "Moral Deskillling and Upskilling in a New Machine Age." *Philosophy & Technology*, 28, 107–124. <https://doi.org/10.1007/s13347-014-0156-9>
- Vallor, S. (2016). *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting*. Oxford: Oxford University Press.
- Vallor, S. (2024). *The AI Mirror: How to Reclaim Our Humanity in an Age of Machine Thinking*. Oxford: Oxford University Press.
- Varela, F. J., Thompson, E., and Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press.
- Verbeek, P.-P. (2011). *Moralizing Technology: Understanding and Designing the Morality of Things*. Chicago: University of Chicago Press.
- Zahavi, D. (2014). *Self and Other: Exploring Subjectivity, Empathy, and Shame*. Oxford: Oxford University Press.

AI Use Statement

AI (Claude, Anthropic) was used to explore and help refine some of the concepts explored in this paper, and in refining the language used. The author takes sole responsibility for the intellectual and scholarly content of the paper.