Check for updates

POSSIBILITY STUDIES & SOCIETY

Theoretical Contribution

Possibility Studies & Society
1–16
© The Author(s) 2026
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/27538699261498153
journals.sagepub.com/home/pst
Sage

The Metaphors We Use, the Possibilities We Lose: AI and the Redistribution of Agency

Danah Henriksen[1] **and Punya Mishra**[1]

Abstract

In this conceptual article, we argue that the metaphors humans use for generative AI are not neutral descriptions; rather, they redistribute creative agency, epistemic authority, and the sense of what is possible between humans and machines, shaping what people can imagine doing with these systems. This redistribution is often absorbed by people rather than chosen: a metaphor becomes part of how someone uses AI without being examined or recognized as a framing. We call this absorbed redistribution and distinguish it from deliberate or resisted redistribution. Drawing on possibility studies, we develop a framework for analyzing how metaphors picture AI, what possibilities they expand and contract, and the conditions under which these effects intensify. Applying it to a set of common AI metaphors, with closer analysis of several that are common in learning, we show that every framing tilts the human-AI relation somewhere, often in different directions. Education offers a clear case, because the metaphors people adopt for AI interact with the metaphors they already hold for learning. We argue that metaphor awareness, recognizing the framings one operates inside and choosing among them, is itself a form of creative and epistemic agency that supports possibility.

Keywords

generative artificial intelligence, metaphor, creative agency, possibility studies, epistemic authority, creativity and education, human-AI interaction, creative thinking

The limits of my language mean the limits of my world.
~ Ludwig Wittgenstein, Tractatus Logico-Philosophicus

A word is dead
When it is said,
Some say.
I say it just
Begins to live
That day.
~ Emily Dickinson

Introduction

Humans live in the world as it is, but we also reach for what might be. We imagine, plan, create, and act toward possibilities that do not yet exist and may never exist—toward the possible (Byrne, 2005). This possibility-seeking capacity emerges in relation to others, mediated by the cultural tools through which we understand ourselves and our context (Bruner, 1990; Glăveanu, 2023). What we can imagine is partly shaped by the tools we have available to imagine with (Glăveanu, 2020). The question becomes which futures we allow ourselves to

[1]Arizona State University, Phoenix, USA

Corresponding author:
Danah Henriksen, Mary Lou Fulton College for Teaching & Learning Innovation, Arizona State University, PO Box 37100, Mail Code 3151, Phoenix, 85069-7100, USA.
Email: danah.henriksen@asu.edu

imagine, and who or what generates those possibilities, *with* us or *for* us.

Generative artificial intelligence (GenAI) is rapidly becoming a technology that we imagine with. In the years since GenAI became widely accessible, it has come to support how people write, learn, ideate, create, decide, and communicate (Pew Research Center, 2025). It has entered classrooms, workplaces, and creative practices, blurring the lines between tool and collaborator, between what humans do and what machines do. Because some forms of AI appear to do things once thought exclusively human, it has sparked debates about the implications for human creativity, agency, and learning (Watkins & Barak-Medina, 2026).

To understand this technology, humans often reach for one of our most fundamental forms of sense-making, that being metaphor (Mishra & Henriksen, 2026). We reach for metaphor out of a need to understand a tool that is both opaque and uncannily fluent, a "black box" that resists inspection yet appears to speak, reason, and create (Bearman & Ajjawi, 2023). AI has been described as a calculator, a collaborator, a genius, a threat, and more (Mitchell, 2024). Metaphors, however, are not merely language, but part of our cognitive infrastructure, and they affect what we notice, expect, and take for granted about unfamiliar domains (Lakoff & Johnson, 1980). The metaphors we choose for AI are not neutral, and their implicit meanings affect what is possible.

AI metaphors extend a long-standing pattern in a new direction: where humans have historically used technologies to describe the mind, we now increasingly also use the language of mind to describe machines, saying that AI "thinks," "learns," "hallucinates," or "creates" (Clark, 2003). Because AI generates fluent language and engages in dialogue, these metaphors can slide into literal assumptions, leading people to attribute human qualities to machines (Cheng et al., 2026). In turn, those assumptions may influence perceived possibilities.

In this conceptual article, we argue that the metaphors used for AI may redistribute creative agency (who is empowered to create), epistemic authority (who decides what counts as knowledge), and perceptions of what is possible between humans and machines. Different metaphors can expand or contract the possibility spaces for people and AI. The critical question is whether this redistribution happens *through* people, by deliberate choice and awareness, or *to* them, as they absorb the metaphors around them, which are often advanced by corporate actors with specific interests (Maas, 2023).

We first situate metaphor within a sociocultural theory of the possible and examine its significance for AI. We then propose an analytic framework for understanding how metaphors expand or constrain possibility within human-AI relations. Using this framework, we analyze the social and educational implications for creativity and possibility thinking when such redistribution occurs without awareness. Finally, we conclude by framing metaphor awareness itself as a form of creative and epistemic agency.

Situating Metaphors as Infrastructures of Possibility

Metaphor is often thought of as figurative language, but it runs much deeper, into the cognitive infrastructure through which we understand the world (Gibbs, 2008). Lakoff and Johnson (1980) showed that human conceptual structures are metaphorical: we make abstract domains clearer by mapping them onto more concrete, familiar ones—a view now well-established across linguistics, cognitive science, and philosophy of mind (Kövecses, 2010). Metaphors structure what we notice in a domain, what counts as relevant, what inferences seem warranted, and what actions seem possible (Draaisma, 2000). For example, if we describe time as a resource to be "spent" or "wasted," the "value" of time comes into focus, while other framings of the concept, such as time being something to inhabit or move through, become harder to formulate. Change the metaphor and you change the thinking.

This matters most when we are trying to understand something new, and when the metaphors available have already been chosen for us (Boucher, 2021).

Familiar things have time to gather their metaphors, which get tested against experience across generations. Novel things lack that history, and the first metaphors used to make sense of them shape later understanding (Ortony, 1993). When early computers were described as "electronic brains," the framing carried into how we now talk about memory, processing, and intelligence in machines, long after the metaphor's limits became visible (Kline, 2015).

AI demonstrates this dynamic vividly, and its contemporary metaphors are wide and varied (Mishra & Henriksen, 2026). It has, for instance, been described as a tool, calculator, search engine, library, assistant, collaborator, tutor, mirror, child, stochastic parrot, oracle, genie, and even God, among others (Vallis et al., 2025). Cheng and colleagues (2025) found that AI metaphors cluster in ways that shape how the technology is perceived, trusted, and used. These dominant clusters have shifted quickly, suggesting that AI's metaphorical landscape is more unstable than that of older technologies. Importantly, the metaphors a community adopts also influence its scientific, legal, and policy responses (Mitchell, 2024), and anthropomorphic language around AI has circulated long enough that it often passes unnoticed, not registering as metaphorical at all (Maas, 2023).

This body of work establishes that AI metaphors are not neutral: they cluster, they shape perception and trust, and they carry consequences for policy and practice. We build on it but make a different move. Rather than asking how metaphors shape perception in general, we ask what they redistribute, and between whom. Our claim is that AI metaphors reallocate specific goods (creative agency, epistemic authority, and the sense of what is possible) along the human-AI relation, tilting each toward the human or the system. And we argue that what matters is not only which metaphor is in play but the mode of that redistribution: whether it happens through the user, who meets the framing with awareness (adopting it deliberately or refusing it outright) or to the user, who absorbs it by default. Naming this, "absorbed redistribution," and grounding it in the field of possibility studies, distinguishes this account from existing work on how AI metaphors shape perception and policy.

Yet, metaphors are not linguistic props to be set aside once the thinking is done. They are how the thinking happens, and they help set the boundaries of what thinking is possible (Ortony, 1993). Possibility studies can offer language for describing not only what metaphors mean but what they allow us to think and do.

The field of possibility studies treats possibility as relational and mediated (Beghetto & Glăveanu, 2022; Hanchett Hanson, 2023). What we can imagine, pursue, or bring into being is not built fresh inside the individual mind. A child imagining their future draws on the lives they have seen, the stories they have heard, and what their world has shown them (Zittoun & Cerchia, 2013). The space of the possible runs through perspective, position, and the cultural and material tools available (Glăveanu, 2020). Bruner (1990) argued that imagination lives in narrative and cultural practice, not simply inside the head. Extending these insights, recent work has brought a possibility lens to dialogue and perspective-taking (Glăveanu, 2023), to the pedagogical conditions that influence possibility thinking (Beghetto, 2023), and to imagination distributed across people and materials (Zittoun & Cerchia, 2013).

Metaphors matter because they do not merely describe, but mediate what people can imagine doing with the thing described (Ortony, 1993). For instance, calling AI a *tutor* brings certain pedagogical futures into focus and pushes others into the background, positioning AI as the holder of knowledge and the human as the recipient. Calling it a *collaborator* foregrounds

questions of shared work, authorship, and responsibility that recede when AI is framed as a tool (Cross, 2025). Calling it an *oracle* minimizes concerns about judgment and verification in ways those other metaphors do not. Each metaphor establishes a different field of possibilities for what humans and machines are doing and redistributes imagined creative agency between them (Kajava & Sawhney, 2023).

Creativity is closely related to possibility, as creative agency cannot be separated from the field of possibility it works within (Glăveanu, 2020). To act creatively is to engage with and extend the space of what we can imagine or make (Hanchett Hanson, 2023). When a metaphor makes a phenomenon thinkable, it does two things. It shapes which futures we can imagine, and it shapes how creative agency gets distributed. With AI, this means the metaphors in circulation are already influencing how creative possibility is distributed between humans and machines, whether or not people recognize it happening.

Historically, misleading metaphors could be corrected over time. By misleading, we do not mean merely selective or partial, as every metaphor foregrounds some aspects and conceals others, but a metaphor that obscures its own limits, making what it excludes difficult for the user to see. Human capacities for metaphor developed where metaphors could be tested against experience, allowing misleading ones to be revised or abandoned. But with AI, that process changes. Humans reach for the vocabulary of the mind to describe what AI does, using words like learning, understanding, intention, care (Kajava & Sawhney, 2023). AI systems are built to invite that vocabulary, offering interfaces and outputs that sound like conversation and reasoning (Mollick, 2024), but these systems do not have minds in the ways the vocabulary suggests. Ordinarily we would notice the mismatch when metaphors rubbed against reality and showed their seams—but AI systems are designed to produce exactly the cues those metaphors predict, so the mismatch is not visible. The seams never show.

The AI Context: The Work That Metaphors Do

For most of technology's history, metaphor flowed in one direction between minds and machines. Humans borrowed from machines to describe the mind, not the other way around. Greek hydraulics gave us the humors, clockwork gave us mental mechanisms, and the digital computer gave us ideas like memory storage, processing, and bandwidth (Draaisma, 2000). We still speak of memory as storage, thinking as processing, and the mind as a mechanism (Epley et al., 2007). And because machines are clearly not minds, the limits of these comparisons have been obvious; we could see that we were borrowing language from one domain to illustrate another (Draaisma, 2000). GenAI adds a new reverse directional flow, however, such that we now also borrow from the human mind to describe an opaque machine, mapping the unexplained onto the unfamiliar. The problem is that the human mind is not well understood, yet our sense of familiarity with our minds can lead us to misread what AI is doing (Vallor, 2024).

This second, opposite-direction flow would matter less if we could hold mind-language at a critical distance, understanding that "learns" for an AI training process does not mean what it means for us. In practice that distance evaporates, mostly because AI reaches us through dialogue, and dialogue is how humans recognize other minds (Epley et al., 2007). Until recently, only people produced fluent, responsive, context-sensitive language, and human social cognition developed under those conditions (Sperber et al., 2010). We are predisposed to attribute minds wherever we encounter mind-like cues in language (Reeves & Nass, 1996). GenAI is distinctive not because it triggers this tendency occasionally but because it does so consistently and at scale. The mind-language we use ("thinking," "understanding") is the trace of a deeper move: treating the system as if it had a mind (Kajava & Sawhney, 2023). Our capacity to recognize persons is now triggered by something that is not a person.

Words like "understands," "learns," or "cares" imply internal states, and in ordinary interaction we do not accept such claims automatically. People constantly, often unconsciously, judge how well such descriptions fit another person (Bielik & Krell, 2025). We judge whether someone truly understands a topic, knows the limits of their knowledge, or is sincerely concerned, based on signals we pick up through interaction (Dennett, 1987). Sperber et al. (2010) call this epistemic vigilance: the cognitive mechanisms that humans use to monitor incoming information for reasons to trust, question, or doubt it.

Vigilance relies on signals from human communication: fluency about complex topics, hedging in the face of uncertainty, hesitation at the edges of one's knowledge, care in phrasing difficult things (Bielik & Krell, 2025). These signals carry information partly because they are costly. Speaking clearly about complex matters requires knowledge and effort. Hedging can indicate self-awareness and restraint. Sustained attention in conversation shows care because a person is spending time and emotional energy rather than directing it elsewhere (Sperber et al., 2010). The cues are imperfect, but they mean something because they are tied to underlying human capacities and commitments.

Contemporary AI systems disrupt this logic. They produce many of the cues that humans evolved to recognize: fluent prose, responsive answers, hedging language when prompted, and the cadence of thoughtful attention. What is missing is the cost behind those cues that normally gives them meaning (Maynard, 2026). AI generates fluent output through computation, without any understanding. Its responsiveness comes from being optimized for interaction, not from attention. Its hedging can reflect training patterns, not self-awareness.

Maynard (2026) characterizes these outputs as "honest non-signals," i.e., cues that accurately reflect system operations yet fail to carry their traditional social meaning. However, applying "honest" to AI risks inadvertently introducing an anthropomorphic framing, implying a relationship to truth or good faith in a system operating entirely outside those bounds. Furthermore, negating the concept as a "non-signal" is misleading since it obscures the fact that a real communicative signal is indeed being transmitted. We instead propose the term *hollow signals:* outputs that maintain the structural surface of human communication (such as conversational fluency, hedging, or deference) while lacking the internal substance of understanding, intentionality, or epistemic risk.

The consequences for AI metaphors are significant. In ordinary life, poorly fitting metaphors get corrected when they conflict with experience (Ortony, 1993). With AI, that corrective friction weakens, because our mind-metaphors predict fluent, responsive, attentive-seeming output, and AI systems produce exactly those signals. But they produce those signals because they are designed to, not because the indicators are real (Bender et al., 2021). The very features that make these systems compelling also help our figurative language drift into literal belief (Cave & Dihal, 2019). Speed compounds the problem. Misleading metaphors are usually revised slowly through debate, counterexamples, and better alternatives (Ortony, 1993), but GenAI evolves so quickly that before one set of metaphors can be examined, another wave of tools has arrived. Users adopt whatever framings are most readily available, often supplied by tech companies, interfaces, and media discourse rather than through slower collective reflection (Maas, 2023).

These weakened checks matter because AI metaphors are not only arising from public discourse—many are actively promoted by actors whose interests they serve. Technology companies, EdTech providers, and policy advocates invest heavily in product design, institutional partnerships, and public messaging that embed particular framings (Kajava & Sawhney, 2023; Teräs et al., 2020). Critical scholarship has noted that the language of "personalized learning" circulates in research and in the

marketing of commercial products, wherein a pedagogical framing and a sales framing become difficult to separate (Pelletier, 2024). The phrase AI tutor did not emerge from collective reflection, rather, it coincided with intense commercial investment. AI-focused education companies drew billions in venture capital in 2025, with investors favoring AI-enabled products (HolonIQ, 2025), much of it flowing to personal tutoring and "copilot" tools. When a framing arrives with significant money and corporate force behind it, the conditions for absorbing it uncritically are stronger. Whether redistribution happens *through* people or *to* people becomes more urgent where persuasive cues are widespread and reflection time is scarce. Thus, shifts in authority and possibility can happen by default rather than by deliberate choice.

An Analytic Frame for Possibility Redistribution

Metaphors and the Distribution of Possibility

Any metaphor used for AI does more than simply describe what AI is. It also assigns roles, positioning the system and the human in relation to one another. Calling AI a "tool" positions the human as the one who acts and the AI as that which is acted upon. Calling it a "collaborator" implies joint authorship, where both parties contribute and can be credited (Cross, 2025). These framings come built into the metaphors by linguistic meanings, foregrounding some features and pushing others into the background (Lakoff & Johnson, 1980).

This is not simply a matter of emphasis. Work on the epistemics of mental models describes the models we reason with as being "ignorance-preserving." By making some inferences available, a model holds others outside what can be considered at all (Arfini et al., 2019), so that the same framing that enhances understanding in one direction leaves ignorance intact in another. A metaphor for AI works this way because it assigns roles. Casting AI as a tutor and the human as recipient does not merely de-emphasize the alternative of a skeptical, verifying human user; it removes the ground on which that alternative would be formulated, so the question of whether to check the system's answers need never arise. What a metaphor contracts, then, involves not only what we are less likely to do but what we are less able to imagine; and a possibility that never surfaces as a question gets redistributed by default rather than by choice.

Although metaphors can be interpreted in different ways, the range of meanings is relatively narrow in practice, and each implies certain roles and behaviors (Gibbs, 2008). When the metaphor positions the system as the one who knows, people are more likely to defer to it, accepting answers uncritically or shifting their own thinking to match the system's phrasing. When AI is understood as something to be used (e.g., the tool framing), judgment stays with the human, and any outputs must be evaluated rather than simply accepted. Whatever the speaker may intend, the metaphor shapes what becomes possible.

This dynamic has a well-known precedent in how learning is theorized. Sfard (1998) suggested that learning itself is structured by two competing metaphors, acquisition and participation, and learning is shaped by whichever metaphor predominates. The same logic runs through our account, as the framing does not simply describe, but structures the possibilities within it. From a possibility studies perspective, metaphors do more than influence how individuals think about AI. They redistribute agency, authority, and possibility across humans and machines. Possibility emerges between people, perspectives, and the cultural tools they must work with (Glăveanu, 2023; Hanchett Hanson, 2023), and metaphors are one such tool. With AI, a metaphor does not only reframe the system; it also repositions the human, since the two are defined in relation to each other. Whatever opens or closes for AI also involves some opening or closing for the human, and vice versa. Adopting a metaphor implicitly suggests who has agency, whose knowledge counts as

authoritative, and what each party can do or imagine (Mitchell, 2024; Vallor, 2024). Our framework makes these implicit suggestions explicit and analyzable.

Two forms of possibility matter most in human-AI interactions. The first is *creative agency:* the capacity to make, originate, contribute, and bring something new into the world (Harris, 2021). Creativity scholars have long argued that creative work is distributed across people, materials, and contexts (Sawyer & Henriksen, 2024), and a possibility perspective shows how that distribution shifts with different metaphors. When AI is framed as a collaborator, creative agency is treated as shared between human and machine, with ideas attributed to the interaction itself and the system credited in part. When AI is framed as a tool, creative agency stays with the human, who is expected to direct, select, and take responsibility for what is produced. The activity may look the same from outside, but the distribution changes what counts as making, who is shaping the work, and what each party contributes.

The second is *epistemic authority:* the capacity to know, judge, and hold warrant for what is true or worth attending to (Goldberg, 2010). Epistemic authority can be distributed in many ways, but metaphors often tilt it toward one party or the other (Sperber et al., 2010). Calling AI an "oracle" places epistemic authority with the system, leaving the human to receive and trust what it delivers, taking an answer at face value or moving ahead without verification. Calling it a "calculator for words" keeps epistemic authority with the human. The system produces outputs, but the human is expected to work with them and judge whether they are correct, useful, or appropriate. This relational picture of agency connects to relational autonomy, which understands autonomy not as individual independence but as a capacity developed and exercised within networks of social and epistemic relations (Marin & Copeland, 2024). Self-trust and critical judgment, then, are cultivated through epistemic relationships rather than achieved in isolation. What a metaphor reshapes, then, is not an individual's independence but a relationally constituted capacity.

Creative agency and epistemic authority are not the only forms possibility takes in human-AI relations. Relational and emotional possibility shift in metaphor-dependent ways too, and we must consider where and how they emerge. The point is structural—metaphors for AI redistribute the human-AI relation along several dimensions, and a possibility frame lets us see this as one phenomenon, not a scattering of effects. It helps to think of metaphors as constraints in the constructive sense (Feiten et al., 2023), as constraints are not merely obstacles that subtract from what is possible but conditions that shape and scaffold the space of creative action in the first place. A metaphor's contractions, in this light, are not simply losses; they give the human-AI relation its usable shape. But this is precisely what raises the stakes of our central question. If constraint is what shapes the possibility space (rather than being in opposition to it), then it is not just important to know what a metaphor constrains but who authors the constraint and whether it is chosen or absorbed. A constraint one adopts deliberately is scaffolding; a constraint absorbed by default is a foreclosure of possibility that one never agreed to.

This lets us be precise about what we mean by a misleading metaphor, a term that might otherwise seem to apply to any metaphor at all, since all of them conceal as they reveal. A metaphor is not misleading merely because it is selective. It becomes misleading when its selectivity turns back on itself: when it obscures its own limits, so the concealment goes unregistered, and when the redistribution it performs is one the user would not endorse if they recognized it. An ordinary metaphor constrains in the open; a misleading one constrains while hiding that it is doing so. The mind-metaphors that AI invites are misleading in exactly this sense. They do not merely emphasize certain features of the system; they present an opaque machine as a

knowing, understanding interlocutor while suppressing cues that would show this mismatch, redistributing epistemic authority toward the system in a way few users would choose deliberately. They shut down options without letting the user see clearly what has been shut down, and, as we argue below, AI weakens the ordinary ways a person might catch this.

A note on how we selected and analyzed these metaphors is in order. We apply this frame to several metaphors used for AI, shown in Table 1. We do not offer this as an empirical survey or a systematic content analysis. The nine metaphors are prominent, recurring framings drawn from public, educational, and scholarly discourse on AI, chosen because they are widely recognizable and because together, they span a useful range, from metaphors that under-attribute capacity to the system (calculator for words, library, stochastic parrot) to those that over-attribute it (tutor, oracle, companion). They are illustrative rather than exhaustive; other metaphors could be analyzed with the same frame. The entries in each column are our analytic readings, produced by applying the framework developed above (creative agency, epistemic authority, and the distribution of possibility) to each metaphor, informed by the literature cited alongside it. We treat a possibility as expanded when a framing makes a role, capacity, or way of relating more available or thinkable, and as contracted when the framing pushes one out of view or makes it harder to formulate. The analysis is thus interpretive and conceptual: its aim is not to measure the effects of metaphors but to demonstrate how the framework makes those effects visible and comparable across cases.

Table 1 presents, for each metaphor, what the framing pictures, what it expands, what it contracts or places beyond imagining, and under what conditions. Later, we examine three of the nine illustrated metaphors (calculator for words, collaborator, and tutor), specifically because they cover a span of redistributive patterns. The table offers breadth; the specific readings offer depth.

Applying the Frame: Metaphors and Their Possibility Effects

Reading across the table, patterns appear that no single metaphor captures on its own. Metaphors that under-attribute agency to AI, like calculator for words, library or search engine, and stochastic parrot, contract what humans might experience the system as capable of, but they preserve human judgment and agency. Metaphors that over-attribute, like tutor, oracle, companion, and threat, contract the human side, ceding creative agency, epistemic authority, or relational capacity to what gets assigned to the system. Metaphors that suggest more symmetrical relations, like collaborator and assistant, do mixed work, expanding and contracting on both sides. The shape of the redistribution varies, but the redistribution itself is consistent across cases. Each framing tilts the relation between human and AI, and the possibility frame lets us see those tilts as connected rather than scattered. Calculator for words, collaborator, and tutor span the range from minimal to maximal anthropomorphism, and each shows a different redistribution pattern.

Calculator for words is among the leanest framings for GenAI. It views the system as a functional device producing outputs from inputs without judgment or stake, and it positions the human in charge of purposes, judgments, and authorship. Users maintain a clear sense of what they are doing with AI and what AI is doing for them, treating outputs as raw material requiring human work to become meaningful. What the metaphor contracts is harder to see, because it flattens AI so thoroughly. Picturing the system as a mechanical function closes off the possibility that AI might offer a thought partner, that its outputs might surprise, or that something at the level of language or pattern might be worth pursuing. Thus, it makes the system narrower than it is.

Collaborator sits at the opposite end, picturing AI as a partner rather than a function. The metaphor expands shared creative agency, distributed authorship, and the imaginative space

Table 1. Metaphor-Mediated Redistribution of Possibility in Human-AI Relations.

Metaphor	What the metaphor pictures	Possibilities it expands	Possibilities it contracts	Conditions
Calculator for words Chiang (2023); Roe et al. (2025)	AI processes language like a calculator processes numbers, producing outputs from inputs without judgment or stake.	Expands… Human roles in judgment, purpose, and authorship; Human clarity about how AI operates/produces outputs; Recognition of AI as functional rather than agentic	Contracts… AI's potential use beyond mechanical function; Humans' openness to being surprised or challenged by system outputs; Human willingness to engage AI as thought partner.	Generative when the framing keeps users alert to their own role. Constraining when the apparent neutrality of "calculation" hides how the system shapes what gets produced.
Library or search engine Vallis et al. (2025); Boucher (2021)	AI holds and retrieves information that humans seek.	Expands… Human access to information across larger scales of content/language; The sense that knowledge held by others is reachable; Human capacity to draw upon and work with what AI organizes	Contracts… Human sense that knowing requires effort, encounter, or relationship; Potential for AI to be engaged as something other than a repository of pre-existing content.	Generative when users treat AI's outputs as starting points for inquiry. Constraining when retrieval is treated as equivalent to understanding.
Stochastic parrot Bender et al. (2021); Mitchell (2024)	AI produces fluent language by recombining patterns from training data without understanding.	Expands… Human recognition of what these systems do mechanistically; Human resistance to overclaims about AI cognition; Clearer footing to demystify the technology.	Contracts… AI's potential to surprise, contribute, or be encountered as anything beyond pattern repetition; Human's openness to creative engagement with the system; The space for something interesting and novel from the system.	Generative as a corrective to anthropomorphic excess. Constraining when treated as the whole truth, since it forecloses inquiry into what users are doing with these systems.
Assistant Mollick (2024); Cheng et al. (2025)	AI follows the human's instructions and carries out tasks on their behalf.	Expands… Human ability to get routine work done with less effort; Human capacity to delegate work to the system; The sense that AI is subordinate to human direction and serves human ends.	Contracts… Human habit of making small judgments themselves; Deeper critical thinking about AI's decisions; Humans' sense that delegation costs something, rather than a neutral convenience.	Generative when the human stays attentive to what is being delegated and why. Constraining when delegation becomes habitual and humans stop noticing or start absorbing relevant details/decisions.
Collaborator Cross (2025); Cheng et al. (2025)	The human and AI work together to produce something neither would produce alone.	Expands… Potential for shared creative agency; Distributed authorship; The imaginative space of collaboration, where ideas emerge from exchange.	Contracts… Sense of distinctness of human voice and ownership; Human confidence that contributions are their own thinking rather than dispersed across human-AI exchange.	Generative when treated as one working mode among many. Constraining when the framing obscures the asymmetries the metaphor cannot honor.

(continued)

Table 1. (continued)

Metaphor	What the metaphor pictures	Possibilities it expands	Possibilities it contracts	Conditions
Tutor Mitchell (2024); Vallis et al. (2025)	AI knows what the human needs to learn and guides them through learning it.	Expands… AI's epistemic authority and perceived pedagogical judgment; Access to instruction at scale, including for learners without other forms of guidance.	Contracts… Human learners' agency over their own inquiry; The social and dialogic dimensions of learning between humans; Human awareness that learning does not require an authority who already knows.	Less constraining where learners hold AI authority critically and institutions support questioning. More constraining where outputs are treated as expert and learning is framed as acquisition.
Companion Cheng et al. (2026); Kouros & Papa (2024)	AI accompanies the human in feeling, thinking, or being, offering presence and support.	Expands… Humans' opportunities to feel connection in moments when other human company is unavailable; Emotional support without managing another's feelings in return; The experience of being heard without judgment, on demand.	Contracts… Human capacity for and access to relationship with other humans; The imagined value of intimacy that requires reciprocity and risk; The slower, vulnerable work knowing another person; Holding solitude without loneliness, where being with oneself is not a problem to solve.	Generative in narrow contexts where the limits of the metaphor stay visible. Constraining when felt connection is mistaken for relationship and human alternatives are eroded by use.
Oracle Vallor (2024); Cave & Dihal (2019)	AI delivers answers to the human, who receives them.	Expands… AI's epistemic authority. Human access to seeming authoritative knowledge across varied questions; Opportunities for humans to have uncertainties resolved quickly; The relief of not carrying questions alone.	Contracts… Human capacity to evaluate, verify, or hold uncertainty; The imagined possibility that not knowing might be valuable; The value of inquiry as a process beyond getting answers; Humans' space to value/exercise their own judgment.	Generative only where users treat answers as provisional and stay in the questioning posture. Constraining where authority is granted and verification is skipped.
Threat or replacement Cave & Dihal (2019); Bostrom (2014)	AI endangers human standing, livelihood, or distinctiveness.	Expands… Human perceptions of AI's substitutive capacity, including in domains long considered human; The imaginative space where AI is taken seriously as a competitor rather than a tool.	Contracts… Humans' sense of distinctiveness, particularly in creative or cognitive work; Humans' confidence in their own worth; Humans' confidence that their contributions are valuable even where AI might do something similar.	Generative when the framing motivates protective attention to human capacities. Constraining when it pre-emptively concedes ground and narrows the human possibility space by anticipation.

of working with another, where ideas emerge from exchange rather than a single mind. The contraction is on the human side. The metaphor obscures the distinctness of human voice and ownership, and when AI is a collaborator, the work becomes harder to see as unmistakably one's own. The collaboration is also asymmetric in ways the metaphor cannot honor, since AI has no stake, accountability, or care; treating it as a collaborator implies it does, leading users to attribute moral weight or relationality that the system cannot carry.

Tutor is the metaphor where the redistributive stakes become most visible, both because it is heavily used in educational discourse and because it interacts directly with deeper metaphors people hold for learning itself. The framing expands AI's epistemic authority, its perceived pedagogical judgment, and human access to instruction at scale, which is substantial for learners without other forms of guidance. But it also contracts learners' agency over inquiry, the social and dialogic dimensions of learning, and the slow work of building understanding through struggle, error, and time. Where learning is framed as acquisition from an authority, the tutor metaphor reinforces and extends that picture to AI; where learning is framed as participation or capacity development through doing, it quietly contracts those alternatives (Sfard, 1998).

These three cases reveal how the same frame produces different redistribution patterns. Calculator for words contracts what AI can be; collaborator contracts human ownership; tutor contracts learner agency, social learning, and developmental struggle. Other metaphors in the table do variations on these moves: companion contracts relational possibilities, oracle contracts epistemic ones, threat contracts confidence in human contribution, and the under-attributing metaphors echo the calculator's logic. In every case, expansion and contraction are relationally and structurally linked: every metaphor constrains, and the constraint is what gives the relation its shape. The question is not whether redistribution happens, but whether it happens with human awareness, i.e., whether the user has chosen the trade-off or absorbed it.

The conditions column points to how each metaphor's redistributive effects are not fixed: they vary depending on how the metaphor is used, the institutional context, and the awareness users bring, noting when a constraint offers scaffolding rather than foreclosure. The redistribution is real, but the sharpness of the effects depends on whether the metaphor is held as a working frame or absorbed as truth. That is the question we turn to next.

Through Us or To Us: The Question of Awareness

The same metaphor plays out differently depending on a person's awareness. A learner who calls AI a tutor but keeps questioning its answers, is less likely to defer to the system than one who has absorbed the framing without reflection. A writer who treats AI as a collaborator faces a subtler challenge: the metaphor can blur whose contribution is whose, and keeping track of one's own creative agency requires actively noticing what came from the human and what came from the exchange. What changes across these cases is not whether redistribution happens but whether it happens through the user, as deliberate choice, or to the user, as absorbed default.

Glăveanu (2020) calls this kind of awareness a meta-position, or the ability to see one's own perspective as one position among others rather than as the way things simply are. Possibility here requires holding more than one perspective at once. A person who cannot see their perspective as a perspective has less room to imagine alternatives or hold onto their own agency. The same logic applies to metaphors. A user stuck inside one metaphor for AI cannot see what it is doing. A user who can step back, see the metaphor as a metaphor, and notice what it includes and excludes can choose how much weight to give it.

Most users do not operate this way. The dominant metaphors for AI arise from a small set of visible actors and circulate faster than they can be examined, so users tend to encounter them as established rather than as choices to make (Maas, 2023). The primary distinction is between redistribution that happens to the user and redistribution that happens through them. We call the first *absorbed redistribution:* when a metaphor, and the shifts in agency, authority, and possibility it carries, becomes part of someone's default way of using AI without being examined, chosen, or even noticed as a framing. Redistribution that happens through the user, by contrast, is redistribution met with

awareness, and it takes two forms. In *deliberate redistribution*, a person adopts a metaphor knowing what it foregrounds and backgrounds. In *resisted redistribution*, a person recognizes a framing and refuses it.

The concept of hollow signals demonstrates precisely how interfaces trigger absorbed redistribution. Consider the common framing of the AI tutor whose "infinite patience" never yields to fatigue or frustration. In human learning, patience and frustration are meaningful precisely because they carry a cost to us: a teacher's patience signals deep investment, while their fatigue signals real stakes, effort, energy spent, and human boundaries. The AI's "patience" costs nothing; it is merely endless computation. Yet, when users absorb this hollow signal, they not only attribute genuine care to the machine, but could also potentially come to see authentic human fatigue as a failure rather than a vital signal of shared effort. The hollow signal may do more than mimic human relationality; over time, it could devalue it. Naming these outputs as hollow signals may help preserve the metacognitive distance that lets users draw on functional fluency without losing sight of human commitment.

Resistance is more than declining a particular metaphor; it is refusing the terms a framing imposes. A writer who senses that calling AI a "collaborator" could blur the ownership of their work may reject the metaphor outright, insisting the system is a tool and the authorship is their own. Thus, they might not choose among framings but still refuse one that feel problematic, in order to avoid contracting their own agency and possibilities. This mode is rarer than the other two, and we treat it more briefly, but it marks where metaphor awareness becomes active rather than merely reflective.

One might expect that users would eventually calibrate, much as audiences learned to read television's framings or to recognize what social media platforms shape for them. But past media metaphors could be revised slowly, through debate, counterexample, and accumulating experience. AI metaphors circulate at product-cycle speed, faster than the collective reflection through which misleading framings might ordinarily be revised.

The Stakes of Absorbed Redistribution

If absorbed redistribution is increasingly common, its effects depend partly on which metaphor is absorbed. Studies of whether AI helps or harms creative thinking, learner autonomy, or self-efficacy report mixed results across contexts and framings (DeSchryver et al., 2025). Our argument helps explain this variation: the same system, encountered through different metaphors, can redistribute agency and authority differently. The technology is not the only variable; the framing is too.

Education is a useful test case, because so much AI use happens in educational settings, and the metaphors people adopt for AI interact with the metaphors they already hold for learning itself. As noted earlier, Sfard's (1998) distinction between acquisition and participation structures how people understand learning itself, and which metaphor predominates shapes what they take learning to be, what teachers do, and what counts as evidence that it has happened. The tutor metaphor for AI fits cleanly inside the acquisition picture. The system has content, the learner needs content, the transfer happens. Within that picture, the worry about absorbed redistribution might look overblown: AI delivers what good teachers deliver, and if it delivers well enough, the learner is served. But this only holds if learning is just about acquisition. If learning is also about participation, struggle, dialogue, and the slow work of becoming someone who can think in a domain (Sfard, 1998), then a tutor metaphor reinforces one picture of learning and contracts the other. Learners absorbing both metaphors may find true participation harder to imagine as a real alternative, not because anyone has argued against it, but because the metaphors in use quietly crowd it out.

The cost here is real. Productive struggle, the experience of working through difficulty without quick resolution, is part of how people develop epistemic and creative capacities (DeSchryver et al., 2025; Sawyer & Henriksen, 2024). When AI as tutor makes answers more available, the conditions that form those capacities can erode without anyone deciding they should. The redistribution happens in what learners think learning is for.

Implications: Awareness and Agency for Deliberate Metaphor Practice

If the effects of absorbed redistribution matter, the question is what to do. The temptation is to argue for better metaphors that support productive human-AI engagement and protect creative agency and epistemic authority from being lost. That is one implication of our argument, to choose our metaphorical language, selecting framings that align with our goals and commitments. But no metaphor for AI is neutral, as every framing expands some possibilities and contracts others. We gain something, we lose something.

The critical point here: if redistribution happens whether anyone notices, then noticing is itself a form of agency. A person who recognizes the metaphor they are operating within, who can step back and see it as one framing among others, has more options than someone operating from absorbed defaults. They can shift metaphors when the situation calls for it, hold multiple framings and use each where it serves, or resist a framing altogether, by refusing the terms a metaphor imposes when those terms would foreclose something they want to protect. None of this requires the perfect metaphor, only allowing awareness to choose among them.

Within education, this might mean working with students and teachers to understand the power of metaphorical framings, asking them to notice the metaphors they are using not as a metalinguistic exercise but as a way of seeing what each framing opens and what it forecloses. A class working with AI could examine what changes when the system is approached as a tutor, a calculator, or a collaborator, and what becomes thinkable or unthinkable under each. The point is not to settle on the right metaphor but to develop the capacity to see that one is always operating inside some framing, and that the framing shapes practice.

This involves meta-cognitive practice with real educational stakes. It treats students not as recipients of instruction about AI but as people whose own framings will shape how they live with these systems for the rest of their lives. Further, this same practice applies outside education, as for instance, a writer who works regularly with AI can develop the habit of asking which metaphor they are operating inside on a given day, and whether that framing serves their voice and goals. A clinician using AI can ask what metaphor the institution has adopted, whether it fits the clinical context, and whether it leaves room for their own judgment. A policy maker drafting regulations can ask which metaphor underwrites the proposed rules and how that shapes outcomes. Metaphor awareness is not the only response to absorbed redistribution; cultural and institutional responses to the actors and design choices shaping dominant framings matter as well. Individual awareness cannot substitute for those responses, but it can preserve agency within the human-AI encounters people navigate now. This is where each person's possibility space gets either narrowed or opened.

Several limits bound this framework. First, we have treated metaphors as if users hold them one at a time, but in practice a single person often operates inside several at once, invoking a tutor in one moment and a calculator the next, or holding framings that pull against each other without resolving them. The redistributive effects we describe are, in lived use, layered and sometimes contradictory. Second, the metaphors we analyze are drawn from Anglophone, largely Western discourse about AI, and which metaphors are available, and what they carry, varies across languages and cultural settings. The logic

of redistribution may be transferable, but its content does not translate cleanly, and a fuller account would need to attend to that variation. Finally, our own framework is not exempt from what it describes. "Redistribution," "possibility space," and the language of expansion and contraction are themselves forms of metaphors foregrounding some aspects of the phenomenon and backgrounding others. The frame cannot step wholly outside the dynamic it analyzes; at best it can hold its own framings at the same critical distance it asks of others. We offer it in that spirit, as one way of making the metaphorical shaping of possibility visible, not as a neutral description of how things are.

Conclusion

We opened by considering not what AI can do, but what futures we allow ourselves to imagine, given the metaphors we have chosen or that have been chosen for us. The argument offers a way of holding that question more clearly: metaphors for AI redistribute creative agency, epistemic authority, and possibility, and what matters is whether the people inside a metaphor have any say in that redistribution or simply absorb it.

When metaphors are absorbed rather than being chosen, they narrow not only what people do with AI, but what researchers and educators think to ask about. That makes this framework a starting point as much as an argument, opening questions about how absorbed redistribution operates across settings from classrooms to clinics, and whether metaphor awareness can be cultivated by design.

The argument points not toward one ideal metaphor, as every framing trades off something, but toward a practice: recognizing the metaphors in use, holding them at a distance, and weighing them critically. This form of creative and epistemic agency grows more important as the conditions for absorbed redistribution intensify. It will not solve AI's broader societal problems, but it keeps the possibility space open for the people doing the noticing. The futures we allow ourselves to imagine exist in this kind of possibility space, and the metaphors we use shape how we approach that space. Examining them, however imperfectly, is how we become the ones doing the imagining rather than the ones being imagined for.

Acknowledgements

The authors used a generative AI tool (Claude, Anthropic) to help examine and refine their own ideas, and for assistance with language editing and organizing the writing. All intellectual contributions, arguments, and conclusions are the authors' own. The authors reviewed and verified all content, including references, and take full responsibility for the final manuscript.

ORCID iDs

Danah Henriksen https://orcid.org/0000-0001-5109-6960
Punya Mishra https://orcid.org/0000-0002-9300-4996

Ethical Considerations

This article is a theoretical contribution and does not report any study involving human participants, human data, human tissue, or animals. Ethical approval was therefore not required.

Author Contributions

Danah Henriksen: Conceptualization; Writing—original draft; Writing—review and editing; Project administration.
Punya Mishra: Conceptualization; Writing—review and editing. Both authors approved the final version of the manuscript.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Declaration of Conflicting Interest

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Data Availability Statement

This is a theoretical article. No new data were generated or analyzed in support of this research.

References

Arfini, S., Casadio, C., & Magnani, L. (2019). Ignorance-preserving mental models: Thought experiments as abductive metaphors. *Foundations of Science*, *24*(2), 391–409. https://doi.org/10.1007/s10699-018-9564-0

Bearman, M., & Ajjawi, R. (2023). Learning to work with the black box: Pedagogy for a world with artificial intelligence. *British Journal of Educational Technology*, *54*(5), 1160–1173. https://doi.org/10.1111/bjet.13337

Beghetto, R. A. (2023). A new horizon for possibility thinking: A conceptual case study of Human × AI collaboration. *Possibility Studies & Society*, *1*(3), 324–341. https://doi.org/10.1177/27538699231160136

Beghetto, R. A., & Glăveanu, V. P. (2022). The beautiful risk of moving toward pedagogies of the possible. In *The Palgrave handbook of transformational giftedness for education* (pp. 23–42). Springer International Publishing.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610–623). https://doi.org/10.1145/3442188.3445922

Bielik, T., & Krell, M. (2025). Developing and evaluating the extended epistemic vigilance framework. *Journal of Research in Science Teaching*, *62*(3), 869–895. https://doi.org/10.1002/tea.21983

Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. Oxford University Press.

Boucher, P. (2021). *What if we chose new metaphors for artificial intelligence?* European Parliament (EPRS PE 690.024).

Bruner, J. (1990). *Acts of meaning*. Harvard University Press.

Byrne, R. M. J. (2005). *The rational imagination: How people create alternatives to reality*. MIT Press.

Cave, S., & Dihal, K. (2019). Hopes and fears for intelligent machines in fiction and reality. *Nature Machine Intelligence*, *1*(2), 74–78. https://doi.org/10.1038/s42256-019-0020-9

Cheng, M., Lee, A. Y., Rapuano, K., Niederhoffer, K., Liebscher, A., & Hancock, J. (2025). From tools to thieves: Measuring and understanding public perceptions of AI through crowdsourced metaphors. *arXiv*. https://arxiv.org/abs/2501.18045

Cheng, M., Lee, A. Y., Rapuano, K., Niederhoffer, K., Liebscher, A., & Hancock, J. (2026). Metaphors of AI indicate that people increasingly perceive AI as warm and human-like. *Communications Psychology*, *4*(1), 8. https://doi.org/10.1038/s44271-025-00376-6

Chiang, T. (2023, February 9). ChatGPT is a blurry JPEG of the web. *The New Yorker*. https://www.newyorker.com/tech/annals-of-technology/chatgpt-is-a-blurry-jpeg-of-the-web

Clark, A. (2003). *Natural-born cyborgs: Minds, technologies, and the future of human intelligence*. Oxford University Press.

Cross, A. (2025). Tool, collaborator, or participant: AI and artistic agency. *British Journal of Aesthetics*, *65*(4), 535–551. https://doi.org/10.1093/aesthj/ayae055

Dennett, D. C. (1987). *The intentional stance*. MIT Press.

DeSchryver, M., Henriksen, D., & Leahy, S. (2025). From friction to synergy: The complex interplay of human creativity and AI. *Possibility Studies & Society*. https://doi.org/10.1177/27538699251350483

Draaisma, D. (2000). *Metaphors of memory: A history of ideas about the mind*. Cambridge University Press.

Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On seeing human: A three-factor theory of anthropomorphism. *Psychological Review*, *114*(4), 864–886. https://doi.org/10.1037/0033-295X.114.4.864

Feiten, T. E., Peck, Z., Holland, K., & Chemero, A. (2023). Constructive constraints: On the role of chance and complexity in artistic creativity. *Possibility Studies & Society*, *1*(3), 311–323. https://doi.org/10.1177/27538699231193539

Gibbs Jr., R. W. (Ed.). (2008). *The Cambridge handbook of metaphor and thought*. Cambridge University Press.

Glăveanu, V. P. (2020). *The possible: A sociocultural theory*. Oxford University Press.

Glăveanu, V. P. (2023). Possibility studies: A manifesto. *Possibility Studies & Society*, *1*(1–2), 3–8. https://doi.org/10.1177/27538699221127580

Goldberg, S. (2010). *Relying on others: An essay in epistemology*. Oxford University Press.

Hanchett Hanson, M. (2023). Reality studies: Reflections on the possible and actual. *Possibility Studies & Society*, *1*(1–2), 87–98. https://doi.org/10.1177/27538699231162311

Harris, D. (2021). *Creative agency*. Palgrave Macmillan.

HolonIQ. (2025). *2026 global education outlook*. https://www.holoniq.com/notes/2026-global-education-outlook

Kajava, K., & Sawhney, N. (2023). Language of algorithms: Agency, metaphors, and deliberations in AI discourses. In *Handbook of critical studies of artificial intelligence* (pp. 224–236). Edward Elgar Publishing.

Kline, R. R. (2015). *The cybernetics moment: Or why we call our age the information age*. JHU Press.

Kouros, T., & Papa, V. (2024). Digital mirrors: AI companions and the self. *Societies*, *14*(10), 200. https://doi.org/10.3390/soc14100200

Kövecses, Z. (2010). *Metaphor: A practical introduction* (2nd ed.). Oxford University Press.

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. University of Chicago Press.

Maas, M. M. (2023, October 25). *AI is like… A literature review of AI metaphors and why they matter for policy* (AI Foundations Report 2). SSRN. https://ssrn.com/abstract=4612468

Marin, L., & Copeland, S. M. (2024). Self-trust and critical thinking online: A relational account. *Social Epistemology*, *38*(6), 696–708. https://doi.org/10.1080/02691728.2022.2151330

Maynard, A. D. (2026). The AI cognitive trojan horse: How large language models may bypass human epistemic vigilance. *arXiv*. https://arxiv.org/abs/2601.07085

Mishra, P., & Henriksen, D. (2026). The mirror and the black box: AI metaphors and what they mean for learning. *TechTrends*, *70*, 559–566. https://doi.org/10.1007/s11528-026-01197-y

Mitchell, M. (2024). The metaphors of artificial intelligence. *Science*, *386*(6723), eadt6140. https://doi.org/10.1126/science.adt6140

Mollick, E. (2024). *Co-intelligence: Living and working with AI*. Portfolio/Penguin.

Ortony, A. (Ed.). (1993). *Metaphor and thought*. Cambridge University Press.

Pelletier, C. (2024). Against personalised learning. *International Journal of Artificial Intelligence in Education*, *34*(1), 111–115. https://doi.org/10.1007/s40593-023-00348-z

Pew Research Center. (2025). How the U.S. Public and AI Experts View Artificial Intelligence. https://www.pewresearch.org/wp-content/uploads/sites/20/2025/04/pi_2025.04.03_us-public-and-ai-experts_report.pdf

Reeves, B., & Nass, C. (1996). *The media equation: How people treat computers, television, and new media like real people*. Cambridge University Press.

Roe, J., Perkins, M., & Furze, L. (2025). Reflecting reality, amplifying bias? Using metaphors to teach critical AI literacy. *Journal of Interactive Media in Education*, *2025*(1). https://doi.org/10.5334/jime.961

Sawyer, R. K., & Henriksen, D. (2024). *Explaining creativity: The science of human innovation*. Oxford University Press.

Sfard, A. (1998). On two metaphors for learning and the dangers of choosing just one. *Educational Researcher*, *27*(2), 4–13. https://doi.org/10.3102/0013189X027002004

Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origgi, G., & Wilson, D. (2010). Epistemic vigilance. *Mind & Language*, *25*(4), 359–393. https://doi.org/10.1111/j.1468-0017.2010.01394.x

Teräs, M., Suoranta, J., Teräs, H., & Curcher, M. (2020). Post-Covid-19 education and education technology 'solutionism': A seller's market. *Postdigital Science and Education*, *2*(3), 863–878. https://doi.org/10.1007/s42438-020-00164-x

Vallis, C., Wilson, S., & Casey, A. (2025). Fear and awe: Making sense of generative AI through metaphor. *Journal of Interactive Media in Education*, *2025*(1). https://doi.org/10.5334/jime.972

Vallor, S. (2024). *The AI mirror: How to reclaim our humanity in an age of machine thinking*. Oxford University Press.

Watkins, R., & Barak-Medina, E. (2026). AI's influence on human creative agency. *Creativity Research Journal*, *38*(1), 137–149. https://doi.org/10.1080/10400419.2024.2437264

Zittoun, T., & Cerchia, F. (2013). Imagination as expansion of experience. *Integrative Psychological and Behavioral Science*, *47*(3), 305–324. https://doi.org/10.1007/s12124-013-9234-2