God, Servant, Master or Friend?
How Metaphors Shape the Way Creative AI is Built and Used

Alison Pease[1], **Margareta Ackerman**[2] and **Dave Murray-Rust**[3]
[1] School of Science and Engineering, University of Dundee, UK
[2] Department of Computer Science and Engineering, Santa Clara University, CA, USA
[3] Faculty of Industrial Design Engineering, Delft University of Technology, Netherlands

Abstract

Creative AI has been characterized through a wide range of metaphors, including a mirror, genie, parrot, collaborator, tool, and magnifying glass, among others. In this work, we bring these metaphors together and organize them along a series of conceptual spectra, considering their implications and the ways they shape our understanding of Creative AI. Such reflection is essential for both interpreting and steering the broader narrative around Creative AI, not only in terms of public perception, but also in how we as a community move forwards in a fast changing socio-technological landscape.

Introduction: The Need for new Metaphors to describe Computational Creativity

At an AI Impact summit in February 2026, OpenAI CEO Sam Altman argued, controversially, that "people talk about how much energy it takes to train an AI model - but it also takes a lot of energy to train a human" (Berger 23rd Feb 2026). In contrast, he previously cautioned against "anthropomorphising" artificial intelligence, emphasising that such systems should be understood as "tools" rather than as "creatures" (Lee 2nd March 2024). In each instance, he employed a different metaphor to suit the point he was making: likening it to human development in the first, to counter concerns about AI's energy consumption; and to tools in the second, to counter efforts to regulate AI. Metaphors for describing creative AI are now deeply embedded in our cultural discourse and are frequently invoked to make a point, often implicitly and without any expectation of consistency or the need to justify the comparison.

Historically, researchers in Computational Creativity have largely tended to limit themselves to long-standing, foundational metaphors such as *agents*, *knowledge*, *goals*, *learning*, *training*, *rewards*, and *artificial intelligence* itself, which have been embedded into the discipline since very early days.[1] Given the relatively small, technically competent audience for our work, communication has usually been between academics themselves and related industries, meaning that clear, technical language and direct, literal communication has been the norm. Of course, outreach activities to the wider public have been held, but generally these are led by people in the field, who frame the narrative themselves.

Since Generative AI has become mainstream, this situation has dramatically changed, with non-experts now routinely talking and writing about Creative AI. As in any context where a complex and not yet widely understood domain enters public discourse, this creates fertile ground for metaphors, and a diverse range has emerged to help to make sense of the new technology. Today. we read about parrots, mirrors, magnifying glasses, hallucinations, drunken postdocs and thieves - depending on which characteristic of AI a narrator currently wishes us to focus upon.

The power of language over thought is well established by cognitive scientists and linguists (Lakoff and Johnson 2008; Thibodeau and Boroditsky 2011), and well understood by those who seek to persuade or explain (Lakoff 2010; Charteris-Black 2006; Gibbs 2008; Sfard 2012). Metaphors do much more than decorate language; they actively shape how we develop, perceive and engage with technologies such as Creative AI, which in turn shapes how they respond to us. Examining the metaphors that frame discussions about Creative AI is crucial for understanding and guiding the narrative surrounding it. This applies not only to public perception, but also to how developers and users envision its appropriate uses and future directions it might take.

We start by delving into previous work on metaphors in generative AI. We then introduce three primary dimensions through which to understand and compare metaphors for creative AI. We then identify common metaphors in CC and consider how they have influenced our field. We conclude by exploring the impact of popular AI metaphors on creative AI, how it is created, used and understood.

Background: The Study of Metaphors for Generative AI

Alongside the emergence of new and colourful metaphors for AI, has emerged a new academic field which studies them. This builds on pre-existing studies of metaphor which have established the key role that they play in our thinking, often without us realising (Thibodeau and Boroditsky 2011). Metaphor hunters often joyfully collect them, pub-

[1] Although these terms have now evolved into technical language, they were originally grounded in metaphorical thinking, much like the Computing Desktop metaphor (Moran and Zhai 2007).

lishing lists and taxonomies which they compare to "picking up seashells on a beach, bringing them home and sorting them into little boxes." (Nerlich 2025).[2] This is now a serious area of study, with metaphorists shifting from simply cataloguing metaphors to critically analysing their social effects.

For example, Brigitte Nerlich examines the lifespan and types of AI metaphors (see also (Bowdle and Gentner 2005)), arguing that as we grow more familiar with creative AI technologies, our metaphors are shifting. Early metaphors emphasised wonder and magic, like "oracle" or "crystal ball" (eg, see (Sharkey and Sharkey 2006)), while later ones adopt more grounded, critical, and pedagogical framings, reflecting a move from asking "What is this thing?" to "How does this technology affect our practices, students, and knowledge systems?" (Nerlich 2025). Nguyen systematically analyses how they're used in the press, specifically the Guardian (Nguyen 2024). Maas examines 55 analogies for AI and explores their policy implications, considering aspects such as what AI is, how it works, how we relate to it, its current or potential uses, and its unintended risks, benefits, or side effects (Maas 2023). Other researchers turn to AI itself for metaphors: Conoscenti has analysed AI-generated metaphors, creating a taxonomy that maps them across various dimensions from the system's own perspective (Conoscenti and others 2025).

Values and Power Structures

As we've seen with Altman's use of different metaphors, we need to give careful consideration to who invents and employs the metaphors, and for what purpose. Metaphors are politically very powerful and we need to ask what are the benefits, risks and harms of the metaphorical framing. Metaphors reflect our values and we should explore them to analyse their implications, and consciously choose ones which align with our values.

Some metaphors have innocent, well-intentioned origins; for instance, the pervasive yet problematic *cancer-as-war* metaphor was formed to drum up support from funding agencies who, it was deemed, were more likely to award funding if the problem was framed in military terminology (Hauser and Schwarz 2015). This was later shown to be problematic (see, for instance, (Malm 2016; Hanahan 2014)), with many arguing that *cancer-as-a-journey* is a more helpful framing. The *AI-as-collaborator* metaphor, while well-intentioned and very fruitful, can also be seen as misleading and potentially harmful. Advait Sarkar argues that describing interactions with AI as "collaboration" obscures the human labour behind these systems, often masking exploitative relationships between AI producers and consumers, particularly affecting workers in the Global South. He suggests that the metaphor *AI-as-tool* is a more accurate and fair representation of its role (Sarkar 2023).

The now common *AI-as-stochastic-parrot* metaphor in (Bender et al. 2021) shows how those in positions of power (in this case Google) try to control the narrative by forbidding the publicising of the metaphor and later firing one of the co-authors after they went ahead and published (News 20th February 2021; Ryan et al. 2024). It is essential that we listen to how minority groups and marginalized populations frame the conversation, as opposed to buying into the metaphors created and used by those in positions of privilege.

Metaphors are used within a context and should be viewed (alongside the people using them) in the context of the work they are meant to do; for instance, deflecting criticism, demystifying processes, stoking passions. In the above examples, Altman has a vested interest in safeguarding his product, while Bender is interested in demystifying the complex data-driven nature of LLM's.

Metaphors for Creative Thinking

Looking at the field of Design and HCI for inspiration, we see metaphors used as part of creative processes - for example Lockton et. al's New Metaphors cards that generate pairings of physical phenomena along with something it could be a metaphor for, for example, "how could a flock of birds be a metaphor for data privacy?" (Lockton et al. 2019). Alves-Oliveira et al. (Alves-Oliveira et al. 2021), used metaphors to re-think robots; "how can a robot as a shoulder parrot be a metaphor for disobedience?" (ibid). Around AI, Murray-Rust et. al pick up a collection of metaphors that can be helpful for designers understanding the technologies, giving alternate framings, and noticing the marginalisation carried out by existing terms (Murray-Rust, Nicenboim, and Lockton 2022; Murray-Rust et al. 2024). Metaphor Shifts is an exercise to find productive and generative ways to relate to the technology for a particular speculative project, using awkward and challenging metaphors to re-think assumptions (Murray-Rust, Lupetti, and Nicenboim 2024). This is further developed as Metaphor Gardening (ibid) where metaphors for technology are cast as things to cultivate, harvest, explore – and to let die when they are out of season. Blythe et. al develop the concept of speculative metaphors: beyond trying to explain technology or persuade people about its meaning, speculative metaphors help us to think into what it might be – the creative side of 'metaphoring back' (Blythe, Lindley, and Murray-Rust 2025).

Three Dimensions of AI Metaphors

This section introduces a small set of dimensions, or conceptual "spectra," that help organize and compare AI metaphors. By positioning metaphors along shared dimensions, we can better understand how different narratives about AI relate to one another and why and how they sometimes conflict with one another. This framework is particularly helpful for understanding popular AI metaphors that profoundly influence popular culture. We note that our aim is not prescriptive - we are not trying to dictate how AI should be viewed - but rather descriptive, seeking to uncover the core dimensions that shape many of today's popular metaphors. In general it is useful to consider metaphors and their purposes together: some are rhetorical, some genuine attempts to understand,

[2] See also Metaphors of AI, by Suhaib Aslam at metaphorsofai.org

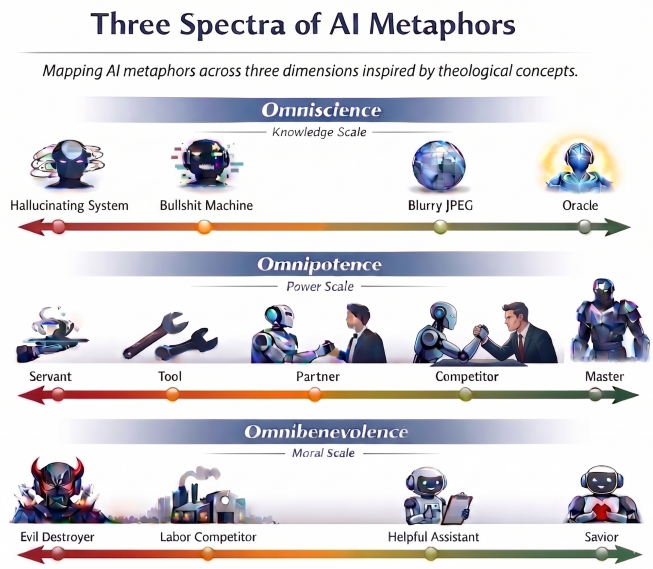

Figure 1: Three primary dimensions through which to understand and compare popular-culture AI metaphor. These dimensions correlate to theological concepts, which have been tied to AI for decades through science fiction literature and film.

some attempts to imagine new possibilities (Blythe, Lindley, and Murray-Rust 2025). For simplicity's sake, in this paper we treat them as distinct from purpose, in order to focus on the characteristics of the metaphor itself.

Our dimensions capture three core traits that people frequently attribute to AI systems: how much they know, how powerful they are, and whether they ultimately benefit or harm society. Interestingly, these dimensions echo longstanding ways humans have reasoned about powerful non-human agents. In theological traditions, similar traits are often discussed through the attributes of omniscience (all-knowing), omnipotence (all-powerful), and omnibenevolence (all-good).

These theological attributes represent extreme ideals. Omniscience implies complete and perfect knowledge, omnipotence implies unlimited power, while omnibenevolence refers to perfect goodness. While in practice, such qualities cannot describe real-world entities, they provide a powerful conceptual vocabulary for imagining what a superhuman agent might be like. Their extremity is part of their appeal: they describe not incremental improvements over human capability, but absolute limits.

The presence of similar ideas in contemporary AI discourse is striking. AI is often described as possessing extraordinary knowledge, vast capabilities, or needing to embody an idealized morality. At the same time, frustration often arises when real systems inevitably fall short of these idealized expectations. In this sense, discussions of AI frequently oscillate between technological reality and aspirational visions shaped by longstanding cultural ideas about superhuman capability.

The three dimensions proposed here are particularly well suited for discussing popular metaphors. Although they do not encompass the full range of metaphors found in academic discourse, we believe that they offer a valuable foundation for the study of AI metaphors.

Omniscience: The first dimension concerns knowledge. In theological discussions, omniscience refers to complete and perfect knowledge, often imagined as encompassing all facts about the past, present, and future. Such knowledge is not just extensive but total, imagining a reality where nothing is unknown or uncertain. This views intelligence through a convergent lens, where questions have clear, often unique, answers. When AI is framed as an oracle or an all-knowing system, the implicit assumption is that knowledge takes this convergent form.

Metaphors surrounding AI diverge dramatically along this axis. At one end of the spectrum, AI is framed as fundamentally unreliable. Concerns center on hallucinations, where models confidently produce incorrect information (we delve deeper into this metaphor below). Related to this is the framing of LLMs as "bullshit machines" - systems capable of generating outputs that may or may not be true, but which lack any intrinsic orientation towards truthfulness or falsity (Hicks, Humphries, and Slater 2024; Gershon 2023).

Moving along the spectrum, another metaphor describes generative models as a "blurry JPEG of the web,"[3] suggesting that they compress large amounts of internet data but reproduce only approximate versions of it

At the opposite extreme, AI is often implicitly framed as an all-knowing oracle: a system believed to possess superior knowledge, possessing greater intellect than the user and perhaps even all of humanity. In this framing, the system becomes an authoritative source of truth, in effect, assumed to possess a form of omniscience.

The influence of this belief is already visible across many domains of everyday AI use. Users defer judgment to AI outputs, incorporating generated text directly into their work, relying on the system's responses without verification, and treating its answers as inherently trustworthy. This behavior contributes to the growing spread of misinformation and the proliferation of low-effort AI-generated content sometimes referred to as "AI slop."[4] In practice, the oracle metaphor encourages users to disengage their own critical reasoning, outsourcing thinking itself to the system. The scale at which such human behaviors are now emerging reveals how deeply influential this metaphor is in shaping how people interact with AI systems today. This dimension captures a central tension in AI discourse: whether generative systems should be understood primarily as unreliable, nonsensical text generators or as new knowledge oracles.

[3] https://www.newyorker.com/tech/annals-of-technology/chatgpt-is-a-blurry-jpeg-of-the-web Accessed: 03/08/2026

[4] https://hbr.org/2025/09/ai-generated-workslop-is-destroying-productivity?utm$_source = chatgpt.com$Accessed: 03/08/2026

Omnipotence: In theology, omnipotence refers to unlimited power: the ability to accomplish anything that is logically possible. An omnipotent being is not constrained by ordinary limits such as strength or resources. In discussions of AI, similar concerns arise around the reach of technological power. Broadly speaking, how much influence can AI systems exert over human activities and economic structures production?

Metaphors along this axis place AI in different positions within a hierarchy of power. At one end, AI is understood as a servant, a tool that executes human instructions and remains firmly under human control. Many traditional software tools fit comfortably within this framing. In the middle of the spectrum, AI may be described as a partner (a framing that is also common in CC literature), contributing capabilities that complement human skills. In creative contexts, this often takes the form of co-creative systems that suggest ideas or assist with exploration.

Further along the spectrum lies the metaphor of the competitor or rival. This position has become especially salient in the current environment, where many worry about the possibility of mass job displacement due to AI, without a corresponding safety net across many societies. At the far extreme lies the metaphor of the master, common in speculative fiction and existential-risk discourse, where AI systems dominate or control humanity.

Omnibenevolence: The third dimension concerns moral orientation. In theology, omnibenevolence refers to perfect goodness: a being whose actions are always aligned with the ultimate well-being of others. Through this lens, the being's actions are understood as good even when harm occurs, assumed to serve a larger purpose surpassing human understanding. When applied to AI discourse, this dimension raises questions about whether increasingly capable systems will ultimately benefit or harm society.

Science fiction frequently explores the negative extreme of this spectrum, portraying AI as an existential threat capable of destroying humanity (ex. The Matrix). These narratives emphasize the risk that powerful systems could act in ways fundamentally misaligned with human interests. At the opposite extreme, AI is imagined as a benevolent technology capable of solving major global challenges, from global warming to freeing people from the burden of work (with a utopian outcome such as the one envisioned by Star Trek). In this framing, AI becomes a transformative force for social good, accelerating scientific discovery, improving healthcare, and addressing complex global problems.

Between these extremes lie more ambiguous possibilities. AI systems may not possess malicious intent yet still produce significant social disruption, for example through labor displacement, misinformation, or the concentration of economic power. This dimension can be represented as going from evil world-destroyer, to a labor competitor, through to assistant and savior. See Figure 1.

While clearly the three dimensions of omniscience, omnipotence, and omnibenevolence are not comprehensive, together they help to map the metaphor landscape of AI discourse. This framing is an import from a theological domain to an epistemic one, which may not come through entirely unscathed; however we see it as a useful construct to see how different metaphors emphasize different positions along these spectra. This helps to explain why narratives about AI can vary so widely and why debates about the technology often appear to talk past one another.

Metaphors used in Computational Creativity

Tool – collaborator – creative entity: By far the most embedded and fruitful metaphorical framing for CC, around which the CC community have long based their work, and indeed the field, is that of *creativity support tool* $\rightarrow$ *co-creative collaborator* $\rightarrow$ *autonomous creative entity*. Discussions have focused around where on this spectrum the focus of CC should lie, with field leaders such as Maher and Grace advocating for the broad space in the middle, and Colton and Cardoso the space at the right hand end. The space at the left hand side, sometimes referred to as 'weak CC' (Veale and Pérez y Pérez 2020), is often dimissed as territory for ('mere' non-CC) tools such as PhotoShop, with CC researchers positioning themselves against these sorts of systems.[5] However there is a small amount of work in CC done at this end as well, such as (Compton and Mateas 2015), as well as work on where the line between tool and collaborator lies, and who decides that (Lawton, Grace, and Ibarrola 2023).

Accompanying research agendas centre respectively around interaction design for *AI-as-co-creative collaborator* and autonomous simulation of aspects of human creativity for *AI-as-autonomous creative entity*. In the former, theoretical concepts focus on aspects of the creative experience, such as conceptual shifts (Karimi et al. 2019) and context-awareness (Ibarrola, Lawton, and Grace 2024), and how co-creative systems can enhance these. Evaluation criteria focus on the quality of the user experience and interactions with domain experts (Kantosalo and Riihiaho 2019); asking questions around who is evaluating the creativity and what is being evaluated (Karimi et al. 2018); comparing various levels of user interaction in different systems (Kantosalo et al. 2020). Key CC systems in this area include Ackerman's interactive songwriting platforms *LyricStudio* and *MelodyStudio* (WaveAI) (Ackerman 2021), which assist users in music composition.

In the latter *AI-as-autonomous creative entity* framing (or 'strong CC' (Veale and Pérez y Pérez 2020)), theoretical concepts centre around the goal of developing automous simulation of aspects of creativity, looking, for instance, at how authenticity (Colton, Pease, and Saunders 2018), framing (Cook et al. 2019), and creative personhood (Pease, Colton, and Banar 2023) might be represented in a computational context. Evaluation criteria in this framing focus on aspects of a system's process or output; such as Ritchie's criteria (Ritchie 2007), Jourdanous's standardised procedure (Jordanous 2012), and the FACE and IDEA models (Colton and Pease 2011). A key CC system here is Colton's *Painting*

[5] In passing, it is worth noting that PhotoShop now has added functionality, including text-to-image capability. As a community, we may need new tools to contrast with 'real' CC systems.

Fool, a long running project which Colton has developed in the hope that one day it would be taken seriously as a creative artist in its own right (Colton 2012). Indeed, while many have attempted to define the field of CC, the definition which has taken hold (as shown by its many citations [6]) lies within this metaphorical framing, in which Colton and Wiggins define CC as developing and thinking about systems which exhibit creative behaviours: "The philosophy, science and engineering of computational systems which, by taking on particular responsibilities, exhibit behaviours that unbiased observers would deem to be creative." (Colton, Wiggins, and others 2012, p1).

Search: Another foundational metaphor is that of *creativity-as-search*, itself relying on the *knowledge-as-space* framing (Newell and Simon 1972). Boden's seminal work (Boden 2004), possibly the single most influential work in the field, relied on these metaphors, with different types of creativity defined by where in a particular conceptual space an artefact or knowledge fragment lay, and how the "knowledge mapper" had got to it. This was a natural fit for AI, as *search* through a problem or knowledge space was a core AI metaphor in symbolic AI from the beginning (Wooldridge 2021) and progress was measured by the efficiency of search techniques, such as breadth/depth/best-first search, greedy search, and so on. Wiggins subsequently provided an algorithmic understanding of Boden's account (Wiggins 2006); and Boden's work has been extensively built upon in the community (see, for instance, (Hung and Choy 2013; Liapis et al. 2021; Ritchie 2006; Ruiz et al. 2015; Farina, Pedrycz, and Lavazza 2024)).

Evolution: *AI-as-evolution* invokes a common metaphor in AI development – genetic algorithms. This is used alongside the *AI-as-search* metaphor, for instance, by Maher in her theoretical work on co-evolutionary design (Maher 2000) and by Browne in his groundbreaking work on using evolutionary algorithms to build systems which can automatically create high quality new games (Browne 2008).

Us: While most work in CC clearly upholds identity boundaries between human and artificial creators, these are starting to blur in some cases, invoking the *AI-as-us* metaphor. For instance, Gill Weinberg gave a keynote presentation at the ICCC in 2017 on his work developing an autonomous robotic drummer, Shimon, which is also a wearable robotic prosthetic for a human amputee drummer (Bretan et al. 2016). Shimon is able to listen to music and respond with different rhythms and behaviours, thus becoming an integral part of the human drummer's physical self and musicianship.

AI Metaphors used in Creative Industries

Much of the public discourse on AI blends together science fiction, futurism, and the commercial systems that are already transforming creative work. Within this broader conversation, certain metaphors have become especially influential in shaping how creative AI is understood. As AI enters creative industries in concrete and economically meaningful ways, these metaphors increasingly shape legal arguments, product design and use, and public perception. In this section, we examine several dominant metaphors in discussions of creative AI: the Thief, the Stochastic Parrot, Hallucinations, the Master, the Muse, and the Partner.

The Thief: One of the most dominant metaphors surrounding creative AI is that of the Thief. Today's commercial generative AI systems are trained on large collections of creative works, such as volumes of musical works on YouTube and Spotify[7] and collections of internet images.[8]

The metaphor is often applied in broad terms. Emphasis is often placed on training data having been taken without permission, and little distinction is made between learning from the corpus and copying a specific protected work. By contrast, when evaluating human creators, there is typically a distinction drawn between inspiration and learning on the one hand, and copyright infringement on the other. Would we ever call a person a "thief" for looking at other people's art and being influenced by it? Within the logic of the Thief metaphor, AI systems are seen as either not truly learning. Consequently, AI outputs are often described as unauthorized reuse rather than as original creative work. This process-based view often assumes denied creativity to the AI regardless of the quality or originality of its output, which are deemed unoriginal nearly axiomatically.

A related metaphor is that of donation rather than theft. It has been suggested that one could think of training data as a form of data donation, similar to blood donation, where many small contributions collectively enable societal benefit[9]. Yet the analogy quickly becomes uneasy. Blood donation is voluntary and informed, while much of the data used to train AI systems was never explicitly offered for this purpose. If AI systems are built from countless human contributions, are we witnessing collective data donation, or something closer to involuntary extraction?

The framing of the Thief, and the involuntary nature of data extraction, has been central in recent litigation. In the music industry, major labels including Universal Music Group, Sony Music, and Warner sued generative music companies such as Suno and Udio, alleging that the companies copied copyrighted sound recordings without authorization[10]. In the text domain, The New York Times sued both OpenAI and Microsoft, alleging that Times content was used without authorization in connection with the training

[6] Cited by 871 on Google Scholar, accessed 8th March, 2026.

[7] https://www.musicbusinessworldwide.com/suno-faces-another-allegation-of-stream-ripping-from-youtube-this-time-in-class-action-led-by-indie-artist/ Accessed: 03/08/26.

[8] https://bmmagazine.co.uk/news/disney-and-universal-sue-ai-firm-midjourney-over-bottomless-pit-of-plagiarism Accessed: 03/08/26.

[9] https://digitalvital.org/ethical-data-donation-public-ai-projects/?utm$_source = chatgpt.com$ Accessed: 03/08/2026

[10] https://www.riaa.com/record-companies-bring-landmark-cases-for-responsible-ai-againstsuno-and-udio-in-boston-and-new-york-federal-courts-respectively/?utm$_source = chatgpt.com$ Accessed:03/08/2026

of LLMs [11]. Some of these cases have resulted in licensing negotiations between AI firms and large rights-holding corporations [12].

Yet, these legal developments expose weaknesses within the metaphor. When compensation is paid, it is typically paid to large rights-holder companies. Individual creators whose works populated the datasets often receive little or no direct benefit. In fact, data set size typically makes it impractical to provide meaningful compensation to every creator in the data set (Ackerman 2025). If corporations are compensated while artists are not, who exactly has been made whole? Who is the injured party, and who is entitled to restitution? The clarity promised by the theft metaphor begins to loosen under close examination.

Certainly, for the large rights holders, the Thief metaphor is strategically powerful. It provides legal leverage for rights-holder corporations and offers a simple moral narrative that resonates publicly. At the same time, its accuracy weakens under closer scrutiny. In addition to confusion around who is the victim of the theft, and who should be compensated for it, this metaphor collapses important distinctions between systems that aim to replicate the style of a specific artist and systems trained to model broad statistical regularities across many creators.

It is also important to note that early generative systems sometimes encouraged imitation. Users were invited to generate content "in the style of" named artists. While work still remains, companies have been making progress towards limiting the ability to use commercial AI systems toward direct imitation of specific artists. It therefore may be argued that while some commercial AI systems started of as thieves, this metaphor has become less applicable over time.

A more precise discussion would separate learning from imitation and would confront the policy implications directly. If training is labeled theft, what follows from that claim? Should the response be prohibition, compulsory licensing, revenue sharing, or structural reform of copyright law? The Thief metaphor mobilizes outrage, but it simplifies a complex technical and economic reality without offering any practical solutions.

The Stochastic Parrot: If the Thief metaphor intensifies moral concern, the Stochastic Parrot[13] attempts to reduce anxiety by positioning AI as a fool, and as such, not a threat. It characterizes AI systems as entities that merely predict the next word or pixel, recombining data without understanding. In creative contexts, this framing often aims to suggest that AI is not truly creative. The lack of understanding is used to delegitimize the machine's capacity for genuine intelligence or creativity and in contrast elevate corresponding human abilities. The implication is that there is little to fear: If AI is only a parrot, then its outputs lack depth. Concerns about displacement of human creativity or human creative labour can then be dismissed or at least minimized.

The derogatory interpretation of the Stochastic Parrot assumes prediction is trivial, and presumes that humans do something that is far more sophisticated. Yet developments in neuroscience, particularly predictive processing frameworks articulated by researchers such as Anil Seth, suggests that human cognition is fundamentally predictive (Seth 2021). The brain constantly anticipates incoming sensory input and updates its internal models. Perception itself can be described as controlled hallucination, the brain's best guess about the world that is refined through feedback. This predictive capacity is far from superficial, but rather central to adaptation and survival.

Creativity also depends on prediction. Humans imagine possibilities, simulate alternatives, and refine ideas through iterative expectation and correction. Both human and machine systems generate by structured guessing. To dismiss AI because it predicts misunderstands prediction as shallow, when it is in fact foundational to intelligence.

There is also a historical irony. For decades, AI research was guided by Turing's Imitation Game, which treated successful imitation of human behavior as evidence of intelligence (Turing 1980). Now that AI systems can approximate human creative outputs across many instances, imitation is framed in a derogatory manner.

Hallucination: Hallucinations have become a deeply engrained part of AI discussions. The metaphor derives from human hallucinations, which can occur as a result of psychedelic drugs such as LSD or psilocybin, or as part of certain psychiatric conditions. Because hallucinations are associated with psychiatric disorders, the term carries strongly negative connotations in everyday language.

It can be argued that hallucinations are necessary for thinking. Perception itself is fundamentally predictive. Neuroscientist Anil Seth (Seth 2021) argues that everyday perception is a form of controlled hallucination. The brain actively constructs reality from sensory input, continuously generating and updating predictions about the world and the causes of its sensations. From this perspective, hallucination is not a failure of the system but a fundamental feature of how minds operate. Imagination can therefore be understood as a natural extension of the brain's capacity to generate internal models of reality.

Our tendency to view full hallucinations negatively reflects a cultural bias. Psychedelics have a long history as part of spiritual practices in many Indigenous communities (Schultes and Hofmann 1992). In recent years, there has been a revival of psychedelic use in mental health treatment, where these substances show significant promise (Garcia-

[11] https://www.courthousenews.com/wp-content/uploads/2023/12/new-york-times-microsoft-open-ai-complaint.pdf Accessed:03/08/2026

[12] https://openai.com/index/axel-springer-partnership Accessed:03/08/2026

[13] The phrase comes from the 2021 paper "On the Dangers of Stochastic Parrots" (Bender et al. 2021). Their goal was to criticize large language models for producing fluent language without true understanding. The metaphor suggests that models statistically remix patterns from training data the way a parrot repeats sounds, except the process is stochastic. While not the intention in the original publication, the phrase has taken on a strongly pejorative tone in popular culture. People often use it to imply that LLMs are mere imitators not meaningfully intelligent systems, that their outputs are empty, meaningless mimicry, and that any apparent reasoning is an illusion.

Romeu and Richards 2018).

In the context of AI, however, the term "hallucination" is used almost exclusively in a negative sense. It refers to situations in which an AI system produces a response that is incorrect, often dramatically so. The mistake may be wildly inaccurate or nonsensical, sometimes with a surprising degree of creativity. Images or videos generated by AI can also be described as hallucinations when they deviate radically from what would normally be expected from the given prompt.

The term hallucination, along with its derogatory tone, makes the most sense when AI is viewed through the lens of the all-knowing oracle. If such a "Wizard of Oz" type entity fails to produce correct output, the result is naturally frustrating (and perhaps even surprising). On the other hand, if we consider AI as a creative machine, which aligns with gen AI's original in generative models, these deviations are neither surprising nor devastating. It becomes entirely natural for AI systems to imagine possibilities and diverge from any single fixed notion of truth.

In creative domains, hallucination becomes a friend rather than a foe. Reclaiming this metaphor may be essential if we want AI to play a meaningful role in the creative future. It encourages us to re-infuse AI with creative capability, move away from the oracle metaphor that currently dominates public discourse, inviting imagination back into both AI systems and our interactions with them. In this sense, the metaphor of hallucination may be one of the most important to rethink if we hope to shape a more creative future for AI.

The Master: At the opposite extreme to the Stochastic Parrot is the Master metaphor. Here, AI is seen as surpassing human creative capacity, so much so that human beings may as well put up the white flag and give up. Artists increasingly question whether there is space for human creators in markets saturated with machine-generated content[14]. The metaphor echoes historical artistic traditions in which apprentices learned under a master. If AI becomes the ultimate master, infinitely productive and creatively superior, what remains for humans?

This view often overlaps with the idea of AI as oracle, a highly capable and seemingly all-knowing entity. In such a vision, human contribution can be seen as inferior or at least unnecessary. The danger of this metaphor is that it encourages premature surrender. If humans assume AI is categorically superior, they may disengage before fully exploring collaboration or differentiation. At present, AI systems are lacking for long form creativity in text and video, tend towards average and commercial stylic choices (thus limiting originality), and are largely outside the scope of inventing new genres.

Zooming within specific domains, there are a wide range of shortcomings for AI generated. For example, there remains a persistent qualitative gap between the images created by AI and those made by professional artists (Rondini et al. 2025), where AI images tend to fall into repetitive patterns (Hintze, Åström, and Schossau 2026). It is far from the master that it may appear to be. The Master metaphor reflects genuine economic anxiety. But by overstating AI's superiority, its harm may become self-fulfilling.

The Muse: The Muse metaphor offers a more constructive frame. Here, AI functions as a source of inspiration rather than replacement. Artists use generative systems to spark ideas by generating unexpected images (Choi and DiPaola 2023), melodic fragments, or lyrical prompts[15] that they then refine and transform. Some creators intentionally use less polished models because their deviations from the norm are inspiring, such as relying on earlier version of DALL-E. This aligns with long-standing artistic practices that rely on chance or cross-domain inspiration.

In many ancient cultures, creative inspiration was understood as a force that came from beyond the self. In Greek tradition, poets often invoked the Muses, divine figures believed to grant inspiration and guide artistic creation (Lamberton 1988). A similar idea appears in Roman thought, where poets described themselves as inspired by a divine inner spirit, a guiding presence associated with creativity and insight (Armstrong 1989).

Computational systems as muses is a well-established concept in computational creativity, with a core CC system in the music domain even named Musebots (Eigenfeldt, Bown, and Carey 2015). The idea that a generative system can be utilized to inspire a person as part of a co-creativity process is foundational idea in the field. Recent work in computational creativity has also applied the muse lens to studying current AI systems (Richards et al. 2025).

At the same time, the metaphor carries important gendered implications. Historically, the muse has often been framed as a passive, feminized source of inspiration rather than an autonomous creator with agency or voice. Feminist critiques therefore emphasize the need to move "from muse to maker", rejecting objectification and reclaiming narrative control so that creative subjects become active producers rather than silent inspirations for others (Gane 2024; Sharma 2022). Despite these tensions, the metaphor preserves a meaningful distinction in creative agency: the AI suggests or inspires, and the human makes and decides. While it does not capture the full scope of AI's capabilities or economic impact, it supports product design choices that prioritize inspiration, exploration and co-creativity rather than full automation. It also offers a lens through which creatives can approach current commercial AI systems, using them to inspire while retaining artistic agency.

The Partner: The Partner metaphor represents a promising direction. In computational creativity, co-creative systems have long been envisioned as collaborators (Lubart 2005). In this framing, AI is capable but does not dominate. The AI may propose, adapt, and respond, while humans retain direction and judgment. Or a different design may be possible, where the AI engages in final decision making. The spectrum of possibilities is at least as wide as it is be-

[14]https://www.theguardian.com/lifeandstyle/2026/mar/06/ai-artificial-intellegence-technology-impact-art-artists Accessed: 03/08/2026

[15]https://lyricstudio.net/

tween human collaborators. Some product developments reflect movement toward this model. Image-generation systems increasingly include editing and refinement tools that shift the emphasis from full automation toward interactive control. A primary example of this movement was when text-to-image models, such as Dall-E and Adobe's Firefly, began to incorporate in-painting, letting users regenerate self-selected segments of an image. In-painting is now a fairly standard feature among AI image generators. Such features do not yet constitute full partnership, but they do signal a shift.

Similarly, generative music platforms, such as Suno, which initially emphasized autonomy have gradually introduced more user control in response to market demand. Many creators do not want replacement, and creators are ultimately a major driving force in creative AI products.

Unlike the Thief or Master, the Partner metaphor does not assume an adversarial relationship. One variation of the Partner model is the Humble Creative Machine (Cassion, Ackerman, and Jordanous 2021; Ackerman 2025). In this framework, the collaborative AI is intentionally designed to preserve human primacy. The system is designed on the assumption that humans have creative capability that have not yet been fully realized. The AI's role is to help users realize that capacity more fully. Such a machine remains flexible across levels of expertise: it can strongly support beginners, while offering lighter touch to advanced creators. Between these poles, it supports a wide spectrum of creative interactions, allowing a person to easily navigate the AI and its use as they grow in their own creative capacity.

Crucially, this form of partnership is structured to foster durable growth. The benefits, changes to a person's creative capability, are retained even after the AI is no longer present. As in an ideal human collaboration, interaction leaves the individual more capable than before. In this model, an AI creative partner becomes a catalyst for lasting creative development.

Concluding Remarks

GenAI has advanced so rapidly now that, even within Computational Creativity, our earlier set of metaphors is no longer sufficient. We need new ways to understand it, and creating a set of metaphors which can help us to engage with it ethically and effectively, and communicate to the wider public, is crucial. Metaphors are neither true nor false, however they are far from neutral; they foreground certain ideas while marginalising others. As a community we need to ask ourselves which ideas we want to highlight and which will not impede understanding if they are concealed.

The distinction between product and process has long been an important one in CC (Jordanous 2016; Colton and Pease 2011). It is worth being mindful of which of these different aspects of creativity the metaphors we have discussed foreground. In fact, most of them emphasise process rather than product: *parrot*, *search*, *thieve*, *evolve* are all verbs showing *how* something is done. *Hallucinations* and *bullshit* could both fall into either category. Interestingly, others, such as *tool*, *collaborator*, *creative entity*, *part of us*, *master*, *muse* and *partner* (as well as *a thief*), don't fall into either category, but metaphorise the type of thing creative AI is, and our relationship to it.

We have discussed three primary dimensions through which to understand and compare metaphors for creative AI, both in CC and in popular culture. The all-knowing framing is interesting: in order for a behaviour to be creative it needs to go beyond typicality and aim for novelty, which in an omniscient context may not be possible. Therefore, in this dimension we may be aiming for somewhere left of all-knowing. A particularly interesting area is towards the left hand side, where hallucinations may give rise to serendipity - a key component of creativity. The all-powerful dimension has been a core tenant for those of the CC community who prioritise automous creative behaviour, while those who prioritise creative interactions may aim to the middle ground of this dimension. The good versus evil metaphor is now more relevant than ever. Previously, our work was extremely niche; while interesting to us, it usually had relatively low impact outwith our community. This situation is now completely different, with creative technologies now threatening the livelihoods of creative professionals and possibly even our ability as humans to be creative.

While we have here, to a certain extent, focused on metaphors in isolation from their context, further work should focus on examining metaphors as they are used "in the wild", alongside the the work that they are intended to do. We should also be mindful of which AI technology a particular metaphor is designed to represent, avoiding the tendency to reduce AI to a single monolithic concept. Careful and intentional examination of our use of metaphors is essential for promoting understanding, inspiring new visions, and shaping valuable creative technologies.

Acknowledgments

We are grateful to our anonymous reviewers for their thoughtful comments, which helped to improve the paper.

References

[Ackerman 2021] Ackerman, M. 2021. WaveAI. wave-ai.net.

[Ackerman 2025] Ackerman, M. 2025. *Creative Machines: AI, Art & Us*. John Wiley & Sons.

[Alves-Oliveira et al. 2021] Alves-Oliveira, P.; Lupetti, M. L.; Luria, M.; Löffler, D.; Gamboa, M.; Albaugh, L.; Kamino, W.; K. Ostrowski, A.; Puljiz, D.; Reynolds-Cuéllar, P.; Scheunemann, M.; Suguitan, M.; and Lockton, D. 2021. Collection of metaphors for human-robot interaction. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference*, DIS '21, 1366–1379. New York, NY, USA: Association for Computing Machinery.

[Armstrong 1989] Armstrong, D. 1989. *Horace*. Yale University Press.

[Bender et al. 2021] Bender, E. M.; Gebru, T.; McMillan-Major, A.; and 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*, 610–623.

[Berger 23rd Feb 2026] Berger, E. 23rd Feb, 2026. Sam Altman defends AI's energy toll by saying it also takes a lot to 'train a human'. theguardian.com/technology/2026/feb/23/sam-altman-openai-energy-use-datacenters.

[Blythe, Lindley, and Murray-Rust 2025] Blythe, M.; Lindley, S.; and Murray-Rust, D. 2025. Artificial intelligence and other speculative metaphors. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference*, DIS '25, 347–356. New York, NY, USA: Association for Computing Machinery.

[Boden 2004] Boden, M. A. 2004. *The creative mind: Myths and mechanisms*. Routledge.

[Bowdle and Gentner 2005] Bowdle, B. F., and Gentner, D. 2005. The career of metaphor. *Psychological review* 112(1):193.

[Bretan et al. 2016] Bretan, M.; Gopinath, D.; Mullins, P.; and Weinberg, G. 2016. A robotic prosthesis for an amputee drummer. *arXiv preprint arXiv:1612.04391*.

[Browne 2008] Browne, C. 2008. Automatic generation and evaluation of recombination games. *PhD thesis. Accessed at https://eprints.qut.edu.au/17025/1/Cameron_Browne_Thesis.pdf*.

[Cassion, Ackerman, and Jordanous 2021] Cassion, C.; Ackerman, M.; and Jordanous, A. 2021. The humble creative machine.

[Charteris-Black 2006] Charteris-Black, J. 2006. *Politicians and Rhetoric. The Persuasive Power of Metaphor*. Basingstoke: Palgrave MacMillan.

[Choi and DiPaola 2023] Choi, S. K., and DiPaola, S. 2023. AI as other: An art-as-research approach to generative AI art practice. In *ICCC*, 357–361.

[Colton and Pease 2011] Colton, Simon, J. W. C., and Pease, A. 2011. Computational creativity theory: The face and idea descriptive models. In *2nd International Conference on Computational Creativity*. Association for Computational Creativity.

[Colton, Pease, and Saunders 2018] Colton, S.; Pease, A.; and Saunders, R. 2018. Issues of authenticity in autonomously creative systems. In *International Conference on Computational Creativity 2018*, 272–279. Association for Computational Creativity (ACC).

[Colton, Wiggins, and others 2012] Colton, S.; Wiggins, G. A.; et al. 2012. Computational creativity: The final frontier? In *Ecai*, volume 12, 21–26. Montpelier.

[Colton 2012] Colton, S. 2012. The painting fool: Stories from building an automated painter. *Computers and creativity* 3–38.

[Compton and Mateas 2015] Compton, K., and Mateas, M. 2015. Casual creators. In *ICCC*, 228–235.

[Conoscenti and others 2025] Conoscenti, M., et al. 2025. Portrait of the AI as a young metaphorist. *QUADERNI DEL CIRM* 6:13–48.

[Cook et al. 2019] Cook, M.; Colton, S.; Pease, A.; and Llano, M. T. 2019. Framing in computational creativity-a survey and taxonomy. In *International Conference on Computational Creativity 2019*, 156–163. Association for Computational Creativity (ACC).

[Eigenfeldt, Bown, and Carey 2015] Eigenfeldt, A.; Bown, O.; and Carey, B. 2015. Collaborative composition with creative systems: Reflections on the first musebot ensemble. In *ICCC*, 134–141.

[Farina, Pedrycz, and Lavazza 2024] Farina, M.; Pedrycz, W.; and Lavazza, A. 2024. Towards a mixed human–machine creativity. *Journal of Cultural Cognitive Science* 8(2):151–165.

[Gane 2024] Gane, M. 2024. Myth of the'muse': a history of the gender roles conceptualised in art. *Plurality* 1.

[Garcia-Romeu and Richards 2018] Garcia-Romeu, A., and Richards, W. A. 2018. Current perspectives on psychedelic therapy: use of serotonergic hallucinogens in clinical interventions. *International Review of Psychiatry* 30(4):291–316.

[Gershon 2023] Gershon, I. 2023. Bullshit genres: What to watch for when studying the new actant chatgpt and its siblings. *Suomen Antropologi: Journal of the Finnish Anthropological Society* 47(3):115–131.

[Gibbs 2008] Gibbs, R. e. 2008. *The Cambridge Handbook of Metaphor and Thought*. Cambridge University Press.

[Hanahan 2014] Hanahan, D. 2014. Rethinking the war on cancer. *The Lancet* 383(9916):558–563.

[Hauser and Schwarz 2015] Hauser, D. J., and Schwarz, N. 2015. The war on prevention: Bellicose cancer metaphors hurt (some) prevention intentions. *Personality and Social Psychology Bulletin* 41(1):66–77.

[Hicks, Humphries, and Slater 2024] Hicks, M. T.; Humphries, J.; and Slater, J. 2024. Chatgpt is bullshit. *Ethics and information technology* 26(2):1–10.

[Hintze, Åström, and Schossau 2026] Hintze, A.; Åström, F. P.; and Schossau, J. 2026. Autonomous language-image generation loops converge to generic visual motifs. *Patterns* 7(1).

[Hung and Choy 2013] Hung, E. C., and Choy, C. S. 2013. Conceptual recombination: a method for producing exploratory and transformational creativity in creative works. *Knowledge-Based Systems* 53:1–12.

[Ibarrola, Lawton, and Grace 2024] Ibarrola, F.; Lawton, T.; and Grace, K. 2024. A collaborative, interactive and context-aware drawing agent for co-creative design. *IEEE Transactions on Visualization and Computer Graphics* 30(8):5525–5537.

[Jordanous 2012] Jordanous, A. 2012. A standardised procedure for evaluating creative systems: Computational creativity evaluation based on what it is to be creative. *Cognitive Computation* 4(3):246–279.

[Jordanous 2016] Jordanous, A. 2016. Four pppperspectives on computational creativity in theory and in practice. *Connection Science* 28(2):194–216.

[Kantosalo and Riihiaho 2019] Kantosalo, A., and Riihiaho, S. 2019. Quantifying co-creative writing experiences. *Digital Creativity* 30(1):23–38.

[Kantosalo et al. 2020] Kantosalo, A.; Ravikumar, P. T.; Grace, K.; and Takala, T. 2020. Modalities, styles and strategies: An interaction framework for human-computer co-creativity. In *ICCC*, 57–64.

[Karimi et al. 2018] Karimi, P.; Grace, K.; Maher, M. L.; and Davis, N. 2018. Evaluating creativity in computational co-creative systems. *arXiv preprint arXiv:1807.09886*.

[Karimi et al. 2019] Karimi, P.; Maher, M. L.; Davis, N.; and Grace, K. 2019. Deep learning in a computational model for conceptual shifts in a co-creative design system. *arXiv preprint arXiv:1906.10188*.

[Lakoff and Johnson 2008] Lakoff, G., and Johnson, M. 2008. *Metaphors we live by*. University of Chicago press.

[Lakoff 2010] Lakoff, G. 2010. Why it matters how we frame the environment. *Environmental communication* 4(1):70–81.

[Lamberton 1988] Lamberton, R. 1988. *Hesiod*. Yale University Press.

[Lawton, Grace, and Ibarrola 2023] Lawton, T.; Grace, K.; and Ibarrola, F. J. 2023. When is a tool a tool? user perceptions of system agency in human–AI co-creative drawing. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*, 1978–1996.

[Lee 2nd March 2024] Lee, L. 2nd March, 2024. OpenAI's sam altman says ai is a tool, not a 'creature'. businessinsider.com/openai-sam-altman-ai-is-a-tool-not-a-creature-2024-3.

[Liapis et al. 2021] Liapis, A.; Martínez, H. P.; Togelius, J.; and Yannakakis, G. N. 2021. Transforming exploratory creativity with delenox. *arXiv preprint arXiv:2103.11715*.

[Lockton et al. 2019] Lockton, D.; Singh, D.; Sabnis, S.; Chou, M.; Foley, S.; and Pantoja, A. 2019. New metaphors: A workshop method for generating ideas and reframing problems in design and beyond. In *Proceedings of the 2019 Conference on Creativity and Cognition*, CC '19, 319–332. New York, NY, USA: Association for Computing Machinery.

[Lubart 2005] Lubart, T. 2005. How can computers be partners in the creative process: classification and commentary on the special issue. *International journal of human-computer studies* 63(4-5):365–369.

[Maas 2023] Maas, M. M. 2023. AI is like… a literature review of AI metaphors and why they matter for policy. *AI Foundations Report* 2.

[Maher 2000] Maher, M. L. 2000. A model of co-evolutionary design. *Engineering with computers* 16(3):195–208.

[Malm 2016] Malm, H. 2016. Military metaphors and their contribution to the problems of overdiagnosis and overtreatment in the "war" against cancer. *The American Journal of Bioethics* 16(10):19–21.

[Moran and Zhai 2007] Moran, T. P., and Zhai, S. 2007. Beyond the desktop metaphor in seven dimensions. *Beyond the Desktop Metaphor—Designing Integrated Digital Work Environments* 335–354.

[Murray-Rust et al. 2024] Murray-Rust, D.; Lupetti, M. L.; Nicenboim, I.; and van der Hoog, W. 2024. Grasping AI: experiential exercises for designers. *AI & Soc* 39:2891–2911.

[Murray-Rust, Lupetti, and Nicenboim 2024] Murray-Rust, D.; Lupetti, M. L.; and Nicenboim, I. 2024. Metaphor gardening: Experiential engagements for designing AI interactions. In *DRS2024*.

[Murray-Rust, Nicenboim, and Lockton 2022] Murray-Rust, D.; Nicenboim, I.; and Lockton, D. 2022. Metaphors for designers working with AI. In *DRS Biennial Conference Series*. DRS2022: Bilbao.

[Nerlich 2025] Nerlich, B. 2025. Metaphors for AI: Three blog posts and a summary. makingsciencepublic.com/2025/11/30/metaphors-for-ai-three-blog-posts-and-a-summary.

[Newell and Simon 1972] Newell, A., and Simon, H. A. 1972. *Human problem solving*. NJ: Englewood Cliffs.

[News 20th February 2021] News, B. 20th February, 2021. Margaret mitchell: Google fires AI ethics founder. bbc.com/news/technology-56135817.

[Nguyen 2024] Nguyen, T. M. 2024. Conceptual metaphors of artificial intelligence and AI development in the guardian newspaper. *VNU Journal of Foreign Studies* 40(4):128–141.

[Pease, Colton, and Banar 2023] Pease, A.; Colton, S.; and Banar, B. 2023. On the notion of creative personhood. In *14th International Conference on Computational Creativity*, 117–121. Association for Computational Creativity.

[Richards et al. 2025] Richards, J. T.; Martino, J.; Bellamy, R. K.; and Muller, M. 2025. Musings on AI muses: Support for human creativity. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems Creative AI Track: Humanity*.

[Ritchie 2006] Ritchie, G. 2006. The transformational creativity hypothesis. *New Generation Computing* 24(3):241–266.

[Ritchie 2007] Ritchie, G. 2007. Some empirical criteria for attributing creativity to a computer program. *Minds and Machines* 17(1):67–99.

[Rondini et al. 2025] Rondini, S.; Alvarez-Martin, C.; Angermair-Barkai, P.; Penacchio, O.; Paz, M.; Pelowski, M.; Dediu, D.; Rodriguez-Fornells, A.; and Cerda-Company, X. 2025. Stable diffusion models reveal a persisting human and AI gap in visual creativity. *arXiv preprint arXiv:2511.16814*.

[Ruiz et al. 2015] Ruiz, F. J.; Raya, C.; Samà, A.; and Agell, N. 2015. A transformational creativity tool to support chocolate designers. *Pattern Recognition Letters* 67:75–80.

[Ryan et al. 2024] Ryan, M.; Christodoulou, E.; Antoniou, J.; and Iordanou, K. 2024. An AI ethics 'david and goliath': value conflicts between large tech companies and their employees. *AI & SOCIETY* 39(2):557–572.

[Sarkar 2023] Sarkar, A. 2023. Enough with "human-AI collaboration". In *Extended abstracts of the 2023 CHI conference on human factors in computing systems*, 1–8.

[Schultes and Hofmann 1992] Schultes, R. E., and Hofmann, A. 1992. Plants of the gods: their sacred, healing, and hallucinogenic powers.

[Seth 2021] Seth, A. 2021. *Being you: A new science of consciousness*. Penguin.

[Sfard 2012] Sfard, A. 2012. Metaphors in education. In *Educational theories, cultures and learning*. Routledge. 39–49.

[Sharkey and Sharkey 2006] Sharkey, N., and Sharkey, A. 2006. Artificial intelligence and natural magic. *Artificial Intelligence Review* 25(1):9–19.

[Sharma 2022] Sharma, N. 2022. From muse to maker: How women in art and literature reclaim their narrative voice. *International Journal of English and Studies* 4(9):114–121.

[Thibodeau and Boroditsky 2011] Thibodeau, P. H., and Boroditsky, L. 2011. Metaphors we think with: The role of metaphor in reasoning. *PloS one* 6(2):e16782.

[Turing 1980] Turing, A. M. 1980. Computing machinery and intelligence. *Creative Computing* 6(1):44–53.

[Veale and Pérez y Pérez 2020] Veale, T., and Pérez y Pérez, R. 2020. Leaps and bounds: An introduction to the field of computational creativity. *New Generation Computing* 38(4):551–563.

[Wiggins 2006] Wiggins, G. A. 2006. Searching for computational creativity. *New Generation Computing* 24(3):209–222.

[Wooldridge 2021] Wooldridge, M. 2021. *A brief history of artificial intelligence: what it is, where we are, and where we are going*. Macmillan+ ORM.