Figurative discourse about AI: A Combined Input Hypothesis analysis

Ghsoon Reda

Public discourse about Artificial Intelligence (AI) relies on figurative language as a semiotic resource for conceptualising the technology. While previous research has examined metaphor and metonymy in such discourse, less attention has been paid to how these processes interact in shaping figurative construals of AI. Drawing on the Combined Input Hypothesis, this study analyses naturally occurring examples from published blogs and articles as instantiations of two broad types of figurative framings: those positioning the technology in relation to human activity and interaction, and those emerging from debates about its capabilities, risks, and limitations. The findings show that metonymy and metaphor perform distinct but complementary functions. Metonymy foregrounds contextually relevant aspects of AI, while metaphor organises these selected aspects into coherent interpretive configurations, shaping how the technology is construed in terms of its capabilities, relation to human agency, epistemic accessibility, and risks. This interaction gives rise to evaluations ranging from trust and collaboration to scepticism, awe, and fear. The study demonstrates the explanatory value of the Combined Input Hypothesis for understanding how metonymic profiling and metaphorical projection jointly structure figurative discourse about AI.

Keywords: AI discourse, Combined Input Hypothesis, conceptual blending, metaphor, metonymy, scaling

1. Introduction

Artificial Intelligence (AI) discourse refers to the ways in which speakers use language to frame and shape the conceptualisation of AI technologies (Johnson & Verdicchio, 2017). A key feature of this discourse is *scaling*: the construction of different degrees of AI capability, agency, and influence. Such scaling is realised through various figurative processes, including *conceptual metaphors* (Lakoff & Johnson, 1980, 1999) and *conceptual integration* (Fauconnier & Turner, 2002), which enable AI to be construed as limited and controllable or, alternatively, as expansive, autonomous, and powerful. For example, typologies of AI metaphors (see Nerlich, 2025a, 2025b) often describe a movement from tool-based and controlled framings towards more human-like/agentive and system-level conceptualisations, a shift commonly associated with transitions from awe at the perceived capabilities of AI to growing concern about its autonomy and impact (see, for example, Cheng et al., 2025, 2026; Inie, Zukerman & Bender, 2026; Oster, McCaleb & Mishra, 2025; Vallis, Wilson & Casey, 2025).

Starting from the high-agency end of typologies, some scholars have examined specific metaphorical framings that shape cultural perceptions of AI, showing how they operate not merely descriptively but as mechanisms of meaning-making that amplify expectations and perceived capabilities. For example, Natale and Ballatore (2020) conceptualise the "thinking machine" as a powerful technological myth surrounding AI. This myth is rooted in the idea of a machine capable of human-like thought, shaping how its capabilities and potential are understood. The scholars identify three dominant patterns in the construction of the AI myth: (1) the recurrence of analogies and discursive shifts, through which ideas and concepts from other fields are employed to describe the functioning of AI technologies; (2) a rhetorical use of the future, whereby present shortcomings and limitations are imagined as being shortly overcome; and (3) the relevance of controversies around AI claims, which they argue should be understood as an integral part of the discourse surrounding the AI myth.

Similarly, Cave and Dihal (2019) examine broader narrative frameworks through which AI is imagined in Western cultures. They show that narratives about AI are structured by recurring hopes and fears, expressed through dichotomies such as *immortality* versus *inhumanity*, *ease* versus *obsolescence*, *gratification* versus *alienation*, and *dominance* versus *uprising*. These narratives shape how the potential consequences of AI are imagined, influencing expectations

about both its benefits and its risks. Taken together with the "thinking machine" myth identified by Natale and Ballatore (2020), these recurring narratives demonstrate how cultural imaginaries can amplify the perceived scale and autonomy of AI, often independently of its actual capabilities.

Beyond cultural and rhetorical amplification, Matthews and Danesi (2019) show how conceptual blending contributes to meaning-making by integrating features from human and machine domains, allowing them to be extended beyond their usual boundaries and to generate new interpretations of AI. In this process, meaning emerges through both alignment and contrast between domains. For example, when a biological concept such as the brain is mapped onto an AI system, the contrast between organic cognition and computational processes is compressed, enabling human cognitive properties to be projected onto the machine. This projection can give rise to what is described within Fauconnier and Turner's (2002) Conceptual Blending Theory as *emergent structure*, whereby the AI system is interpreted as possessing intelligence or agency that appears to exceed its underlying computational basis. In such cases, AI systems may be perceived as human-like, despite lacking the embodied and semiotic grounding characteristic of human communication. As a result, metaphorical expressions such as *AI thinks* or *AI learns* allow human cognitive attributes to be transferred onto machine representations, thereby expanding perceptions of AI capabilities.

An important mechanism relevant to Matthews and Danesi's (2019) account is *metonymy*, which plays a key role in structuring the conceptual material available for cross-domain mapping. In AI discourse, this can be seen in the way expressions such as *think* or *the brain* provide access to broader cognitive capacities within the human domain. This type of metonymic selection foregrounds particular aspects of cognition, making them available for metaphorical projection onto AI systems. When mapped across domains, these selectively activated features support the attribution of human-like properties, giving rise to interpretations of intelligence or agency in systems that are otherwise computationally bounded. This notion of metonymy is rooted in Langacker's (1986) distinction between a conceptual *base* and the *profile* designated within it. Linguistic expressions evoke a broader conceptual domain while directing attention to a particular substructure. Metonymy exploits this asymmetry by using one salient element of a conceptual domain to provide mental access to another element within the same domain, including (but not limited to) cases in which a part stands for the whole. It therefore involves a within-domain association rather than a cross-domain correspondence (Barcelona, 2015).

A related function of metonymy in AI discourse concerns referential construal rather than conceptual profiling. As shown in Nerlich (2026), specific technologies such as Large Language Models can function as stand-ins for AI as a whole, creating a subcategory-for-category relationship in which one prominent technology comes to evoke a heterogeneous field of systems and applications. This metonymic construal can encourage the generalisation of the capabilities and limitations of a particular technology to AI more broadly, potentially producing overextended or misaligned interpretations of what AI is or what it can do.

Despite significant advances in the study of AI discourse, a key gap remains in the literature. Existing research has examined metaphor, conceptual blending, and metonymy as important mechanisms of figurative meaning making. However, these processes are often treated either in isolation or within Conceptual Blending Theory, where the focus is primarily on how elements from multiple input spaces are projected into a blended space, yielding an emergent structure. While this approach has been highly influential in explaining cross-domain integration, it offers a less explicit account of how metonymy and metaphor interact to shape context-dependent variation in how AI is conceptualised.

In fact, AI is not discussed in terms of a single, stable conceptualisation. The same technology may be framed as *a tool*, *a collaborator*, or *an autonomous agent*, depending on which aspects of AI are made salient in a particular communicative context. Metaphor analysis identifies these alternative conceptualisations by showing the broader interpretive frameworks through which AI is understood. However, it does not explain which particular aspects of AI are foregrounded to make these conceptualisations meaningful. For example, conceptualising AI as a tool presupposes that its instrumental function is foregrounded while its apparent agency is

backgrounded. Metonymy analysis accounts for this selective profiling by revealing how particular aspects of AI become cognitively salient. Examining AI discourse through both metaphor and metonymy therefore provides a more complete account of figurative meaning-making, revealing not only how AI is conceptualised but also what particular conceptualisations emerge in different communicative contexts.

This complementary perspective is consistent with research on metaphor-metonymy interaction. Goossens (1990) introduced the notion of *metaphtonymy* to describe cases in which metaphor and metonymy operate together, while subsequent studies have shown that metonymy can motivate, constrain, or facilitate metaphorical interpretation (Barcelona, 2000). Building on this work, the present study argues that the characteristics of AI foregrounded through metonymic profiling condition the trajectory of metaphorical scaling. For example, when AI is foregrounded through its instrumental function, it is readily construed through the AI IS A TOOL conceptual metaphor (i.e., cross-domain mapping, Lakoff & Johnson 1980, 1999), resulting in a relatively low attribution of agency. When conversational ability or independent decision-making is foregrounded, AI is instead construed as a partner or an autonomous agent, yielding progressively higher levels of metaphorical scaling. Metonymic foregrounding therefore guides the metaphorical construal of AI by determining which characteristics become available for conceptual projection and, consequently, which contextual inferences are licensed. Moreover, because metonymic profiling is itself context-dependent, the interpretation of the same AI metaphor may also vary across discourse contexts, shaping evaluations that range from trust and admiration to anxiety and fear.

Although the phenomenon of scaling in AI discourse has been widely noted, particularly in relation to shifts in perceived agency, autonomy, and risk, it is typically described in terms of representational outcomes rather than analysed as the result of underlying semiotic and cognitive processes. What remains under-theorised is how variation in these representations, both across and within metaphors in different contexts, emerges from the structured interaction between metaphor and metonymy, and how this interaction regulates interpretive scope and conceptual scaling. This study addresses this gap by applying the Combined Input Hypothesis. This hypothesis, developed by Francisco José Ruiz de Mendoza Ibáñez and colleagues (e.g., Ruiz de Mendoza, 1998; Ruiz de Mendoza & Díez, 2002; Ruiz de Mendoza & Peña, 2002), offers an alternative to Fauconnier and Turner's (2002) Conceptual Blending Theory. It is an integrative framework that models conceptual integration not in terms of emergent blended structure, but as the outcome of structured interactions between metaphor and metonymy operating on a single, dynamically organised input configuration. In this process, metonymy functions as a domain-access mechanism that selects and foregrounds relevant aspects of a conceptual domain, thereby defining the scope of the material available for interpretation. Metaphor, in turn, establishes systematic correspondences between selected elements across domains. Their interaction is regulated through a balance of correlation and contrast: correlation enables the alignment of inferentially relevant structure across domains, while contrast constrains projection by foregrounding non-aligned features that limit interpretive extension.

This perspective is particularly useful for the study of AI metaphors, where variation in interpretation does not depend on the introduction of new conceptual structure, but on how existing features of AI, such as agency, autonomy, or control, are selectively foregrounded or backgrounded across contexts. The Combined Input Hypothesis is therefore employed as an analytical framework that allows these interacting operations to be examined systematically, without positing additional emergent structure. Rather than assuming that meaning arises from a separate blended space, the analysis treats interpretation as arising from the interaction of metonymic selection and metaphorical projection within a single configuration. Differences in how AI is understood can thus be traced to how domain-relevant features are selected and projected, giving rise to shifts in perceived agency, capability, and risk. These shifts constitute processes of conceptual scaling in which interpretive scope is expanded or constrained depending on which features are made salient. Such shifts also license evaluative and affective inferences, such as awe, trust, or fear, which emerge as consequences of how AI is construed, rather than as

properties directly encoded in the metaphor itself. The research questions guiding this study are therefore formulated as follows: (a) How do metonymic processes interact with metaphor to regulate the selection and projection of features in AI representations? (b) How does this interaction scale perceptions of AI in terms of agency, intelligence, and risk, and generate associated evaluative and affective inferences?

The paper is structured as follows. Section 2 presents a review of existing research on AI metaphors, focusing on their evolution and the typologies developed to describe them. In Section 3, I introduce the Combined Input Hypothesis (CIH) and discuss its relevance for analysing the interaction between metaphor and metonymy in discourse about AI. Section 4 outlines the methodology, including the principles guiding example selection and the procedures used to identify the metonymic and metaphorical processes involved in the construals analysed, and Section 5 presents the findings, based on authentic examples drawn from published articles and blogs to examine the trajectory of metaphorical construals of AI, the conceptual scaling they exhibit, and the evaluative inferences that arise from them. Section 6 concludes with a summary of the key findings and suggestions for future research.

2. AI metaphors: Evolution and typologies

The development of AI metaphor scholarship reflects a shift from early functional and descriptive accounts to more critical and systemic forms of evaluation and interpretation (see Nerlich, 2025b). Early work focused on aligning metaphors with users' mental models to ensure metaphorical fit, emphasising usability and predictable functionality. As the socio-technical footprint of AI has expanded, scholars have begun to classify metaphors along two interacting dimensions: *functional role* (e.g., tool, social partner, systemic hazard) and *emotional or psychological driver* (e.g., awe, trust, concern). This dual-axis framework enables the systematic mapping of AI perceptions across two distinct scaling trajectories: upward scaling towards increased agency and systemic impact, and downward scaling towards depersonalised, non-cognitive construals that emphasise imitation or reflection.

Cheng et al. (2025) highlight how AI is frequently framed in low-agency, instrumental terms, while also allowing for shifts towards more agentive and potentially adversarial representations. They discuss a range of metaphorical construals, from AI as a tool to more negatively framed entities such as a thief, which point to variation in perceived agency and intent rather than a strictly defined scale. At the lower end of this range, AI is represented as functional machinery, with phrases such as *calculators*, *libraries*, or *toolboxes*, where human competence dominates and machine agency is minimal. Here, trust is grounded in control and predictability, and the dominant orientation is one of utility, with AI understood as a passive extension of human intent. Complementing this, Vallis et al. (2025) foreground the affective dimension of AI representations, tracing how perceptions can shift along a trajectory from awe to fear. Within this perspective, low-agency, tool-based framings tend to attenuate emotional intensity by emphasising functionality and control, whereas more agentive or opaque representations invite heightened responses, including uncertainty, concern, or existential anxiety.

Related pedagogical approaches (e.g., Zirenko et al., 2025) examine how learners use metaphors to conceptualise AI and learning with AI by mapping it onto familiar domains such as machines, cognition, and opacity. Such metaphors primarily function as strategies for managing epistemic uncertainty, allowing users to make sense of systems whose internal processes are not directly accessible. Within broader AI discourse, the *black box* metaphorical expression similarly foregrounds this limitation by emphasising the opacity of complex systems and the difficulty of tracing internal causal mechanisms. While this framing can shift attention from attributed intelligence to observable input-output behaviour, it also leaves interpretive space open, which may either support more cautious, function-oriented understandings or contribute to perceptions of hidden agency. In the latter case, such framings have been critiqued as forms of "enchantment", where opacity is reinterpreted as depth or latent intelligence rather than structural limitation (e.g.,

Campolo & Crawford, 2020). Similarly, human-like framings such as AI as a *partner*, *tutor*, or *friend* increasingly attribute social and cognitive agency to AI systems and have been associated with what Cheng et al. (2026) describe as "vulnerable trust", where users rely on AI despite recognising its fallibility. Here, epistemic uncertainty is partially replaced by relational trust, and AI is construed as an interactive participant rather than a passive system.

Recently, Critical and media discourses have increasingly mobilised high-intensity risk metaphors, particularly in relation to concerns about safety, governance, and control. Expressions such as *AI apocalypse*, *autonomous threat*, and *alien intelligence* are commonly used in public and journalistic discourse and exemplify a shift towards existential framing, in which AI is constructed not as a tool or assistant but as a potentially uncontrollable force. In some formulations, this extends to the idea of a non-human agency capable of "hacking" the cultural operating system of civilisation (Harari, 2023), further intensifying perceptions of systemic penetration and loss of human control. This metaphorical escalation has also been described in terms of evolutionary and successor framings (Bostrom, 2014), and, in Maas's (2023) account, through a broader atlas of analogies in which AI development is embedded within narratives of competition, acceleration, and existential risk (pp. 17-28). A contrasting development is found in critical linguistic accounts such as Bender et al. (2021), who characterise large language models as *stochastic parrots*. In this framing, AI is reconfigured as a system of statistical patterning rather than understanding, and the metaphor functions as a deliberate demystification strategy that reduces perceived agency and cognitive depth. Earlier narratives that attribute human-like intelligence or even awe-inspiring capacities are thereby reinterpreted in more constrained, non-intentional terms.

A related but distinct critical perspective is developed by Vallor (2024), who conceptualises AI as a mirror in *The AI mirror: How to reclaim our humanity in an age of machine thinking*. Rather than emphasising imitation or randomness, Vallor's framing highlights reflexivity: AI systems are understood as reflecting back patterns of human language, values, and biases embedded in their training data. However, this reflection is not neutral. It is shaped by socio-technical conditions that selectively amplify certain aspects of human behaviour while obscuring others, thereby producing a distorted or partial image of humanity itself. In this sense, the *mirror* metaphor not only constrains AI agency but also redirects interpretive attention towards human responsibility and self-understanding, shifting the focus from what AI *is* to what AI *reveals about us*.

The above examples show that scaling in AI metaphors operates bidirectionally, as metaphorical construals both expand and constrain perceived agency, control, and scope depending on the framing used. This variation gives rise to a corresponding range of affective inferences, from reassurance and perceived controllability at the baseline, through calibrated trust and uncertainty in intermediate cases, to awe and fear at the upper end, and finally to scepticism and demystification in downward-scaled representations.

While much of the literature has examined AI metaphors in terms of relatively stable mappings between source and target domains, along the lines of Conceptual Metaphor Theory (Lakoff & Johnson, 1980, 1999), less attention has been paid to how the same metaphorical expression can support different interpretations across contexts. From this perspective, variation arises not from changes in the conceptual metaphor itself, but from shifts in which dimensions of the underlying conceptual domain are foregrounded or backgrounded in use. Take as an example the *stochastic parrot* expression, which brings together features of computational language processing and imitation. Depending on contextual emphasis, the same framing may support interpretations of AI as a limited imitator or a system that simulates understanding, and can in some cases be associated with tool-like, non-agentive construals of AI.

All this suggests that metaphorical meaning in AI discourse is not fixed at the level of domain mapping, but rather emerges from the selective activation of conceptual material that determines what is available for interpretation. It is this selective activation that regulates how AI is scaled in discourse: by foregrounding certain properties (e.g., agency, autonomy, communicative capacity) and backgrounding others, metaphors expand or constrain perceived

capabilities across contexts and thereby license evaluative and affective inferences such as awe, trust, uncertainty, or fear.

However, the interaction between selective activation and metaphorical mapping has not yet been fully theorised within a unified account of contextual variation in AI discourse. As stated in Section 1, the Combined Input Hypothesis offers a way to address this gap by providing a principled framework for analysing how figurative language structures the interpretation of AI and how these interpretations vary systematically across contexts, giving rise to patterns of conceptual scaling and affective inference.

3. The Combined Input Hypothesis

The Combined Input Hypothesis (CIH) was developed by Ruiz de Mendoza (1998), Ruiz de Mendoza & Peña Cervel (2002), and Ruiz de Mendoza & Pérez (2003), and further elaborated by Ruiz de Mendoza & Santibáñez Sáenz (2003). While it draws on the structured mappings of Fauconnier and Turner's (2002) Conceptual Blending Theory and the processing constraints of Relevance Theory (Sperber & Wilson, 1995), it differs from both by offering a unified account of figurative meaning grounded in interacting cognitive operations.

Unlike Relevance Theory, which often treats metaphor as a loose use of language that gives rise to implicatures, CIH analyses metaphor and metonymy as mechanisms for the generation of *explicatures*: aspects of the explicitly communicated content of an utterance that are pragmatically enriched beyond its encoded linguistic meaning (Sperber & Wilson, 1995). Within the CIH, figurative meaning forms part of the development of an utterance propositional content. At the same time, the CIH departs from Conceptual Blending Theory by rejecting the need for a separate blended space with emergent structure: all relevant structure is derived from the input domains, and apparent novelty arises through inference over this structured material.

Central to CIH is the notion of a *combined input*, whereby multiple conceptual structures (source and target) are activated simultaneously within a single mental workspace prior to projection. Integration within this space is achieved through two main processes: *schematic enrichment* and *integration by combination*. Schematic enrichment relies on *image schemas* (i.e., recurrent patterns of embodied experience such as CONTAINER, PATH, or FORCE (Lakoff & Johnson, 1999), which provide abstract structural templates for organising conceptual content. These schemas do not supply detailed meaning, but rather a skeletal structure that can be filled with domain-specific information. Integration by combination, on the other hand, recruits additional conceptual material where required for the interpretation to be coherent.

These processes are regulated by cognitive operations such as *correlation* and *contrast*. Correlation underlies metaphorical mapping by linking domains that are experientially or culturally associated (e.g., MORE IS UP), thereby licensing cross-domain mapping and projection. Contrast, on the other hand, highlights differences between domains, enabling more complex mappings that rely on partial alignment rather than simple similarity. In addition, *metonymic expansion* and *metonymic reduction* function as domain access operations: expansion activates a broader conceptual frame from a salient element, while reduction foregrounds a specific aspect by suppressing irrelevant structure.

These operations are guided by the following higher-level processes: *cueing, abstraction, integration, and projection*, which determine how linguistic and contextual information activates and structures conceptual material. Cueing directs attention to relevant domains, abstraction provides the level of generality required for cross-domain mapping, integration organises the selected material into a coherent structure, and projection yields the final interpretation. Importantly, metaphor and metonymy within CIH are understood as interdependent operations: metonymy regulates access to conceptual material, while metaphor projects this material across domains.

4. Methodology

This study adopts a theory-driven qualitative approach that applies the CIH to examine how metaphorical construals of AI emerge through the interaction between metonymic profiling and metaphorical projection. The selected conceptual metaphors represent two broad tendencies in public discourse about AI documented in previous research (see Nerlich, 2025a, 2025b, 2026, for an overview). The first comprises widely circulating conceptualisations that have become established ways of framing AI, including AI IS A TOOL (Cheng et al., 2025), AI IS A CO-PILOT (Woods, 2025), and AI IS A BLACK BOX (Mishra & Henriksen, 2026). The second comprises problematising construals that have gained prominence in debates surrounding AI capabilities, risks, and limitations, including AI IS AN ALIEN INTELLIGENCE (Mitchell, 2025), AI IS AN EXISTENTIAL THREAT (Bostrom, 2014), and AI IS A STOCHASTIC PARROT (Bender et al., 2021). The former framings are more likely to function as non-deliberate metaphorical framings in everyday discourse because they draw on familiar ways of talking about AI without requiring conscious attention to the metaphorical framing, whereas the latter are more likely to involve deliberate reflection on how AI is conceptualised through explicit consideration of what AI is and how it can be understood or characterised. However, this distinction is not absolute, as established metaphors may also become objects of deliberate reflection when their adequacy as representations of AI is explicitly questioned, as illustrated by Lawrence's (2025) *Stop treating AI like a tool*. Examining both established and problematising metaphorical framings enables the analysis to capture how different figurative construals foreground different dimensions of AI, including capability, agency, epistemic accessibility, and potential consequences.

Rather than representing isolated linguistic choices, these conceptual metaphors occupy different positions along a continuum of conceptual scaling, ranging from construals that minimise AI's attributed agency, capability, and autonomy by foregrounding human control, to construals that progressively amplify its perceived cognitive capacity, epistemic distance, and potential risks to humanity. AI IS A STOCHASTIC PARROT complicates this continuum by challenging attributions of intelligence or intentionality despite acknowledging sophisticated linguistic performance. The examples analysed in the following section are authentic instances of these metaphorical framings collected from published blogs and articles. They were identified through targeted web searches using the relevant metaphorical expressions and related formulations as search terms. Examples were selected as representative instances of naturally occurring uses of these framings in contemporary public discourse about AI.

The identification of the metaphorical and metonymic dimensions of the examples was guided by established principles in cognitive linguistics. Metaphors were identified where AI was construed through systematic cross-domain mappings in which a source domain provided the conceptual structure for interpreting the target domain (Lakoff & Johnson, 1980, 1999). Following Croft (2008), these mappings were analysed as establishing *iconic* (i.e. resemblance-based) correspondences between source and target domains, while also recognising, in line with Zlatev and Moskaluk (2022), that such correspondences are motivated and interpreted within specific communicative contexts rather than reflecting objective similarities.

Metonymies were identified where one salient aspect of the AI domain provided mental access to another aspect of the same domain, reflecting a within-domain mapping motivated by conceptual contiguity (Langacker, 1986; Barcelona, 2015). The analysis first identified the aspect of AI that was made cognitively prominent in the discourse. It then examined whether this profiled element provided mental access to the broader concept of AI or to another closely related aspect of the same conceptual structure through conceptual contiguity. Expressions were therefore analysed as metonymic when they selectively foregrounded one dimension of AI while backgrounding others, thereby constraining the conceptual representation that subsequently became available for metaphorical construal.

5. Findings

The CIH analysis implemented in this section systematically explains how metonymic foregrounding and metaphorical projection interact in shaping conceptual scaling and the affective inferences associated with different figurative construals of AI. To make this progression explicit, the examples are presented as a trajectory across different dimensions of conceptualisation, beginning with metaphors that construe AI in terms of instrumentality and human-directed functionality, and extending towards construals that foreground broader forms of capability, cognitive difference, epistemic limitation, or potential consequences. The final example, AI IS A STOCHASTIC PARROT, departs from this trajectory by critically re-evaluating assumptions about intelligence and understanding, while illustrating how the same metaphorical framing can support alternative interpretations through shifts in metonymic profiling. This ordering demonstrates that scaling is not a strictly linear process but a dynamic one in which figurative construals may amplify, constrain, or reconfigure different dimensions of AI conceptualisation depending on the interaction between metonymic profiling and metaphorical projection.

5.1. AI IS A TOOL

The AI IS A TOOL metaphor represents a low level of conceptual scaling, whereby the conceptual category of tools is extended beyond its prototypical domain of physical implements to structure the interpretation of AI in terms of instrumentality, limited functionality, and human-directed use. Through this construal, AI is primarily understood as an extension of human capabilities, while capacities associated with independent action and autonomous decision-making are backgrounded. Examples (1) and (2) illustrate how this metaphor is instantiated in naturally occurring discourse.

(1) *AI is a powerful tool that can enhance your writing, but don't rely on it completely.* (Sira, 2024)[1]
(2) *Practice using AI regularly to improve your skills and become a master wordsmith.* (Sira, 2024)

In these examples, the tool metaphor is reinforced not only through the explicit lexical cue *tool*, but also through contextual information that positions AI within a human-directed activity. In Example (1), the additional warning not to rely on AI completely further reinforces this constrained construal by presenting the technology as a supportive resource whose use remains subject to human judgement and control. Example (2) similarly frames AI as a means through which users develop their own abilities, maintaining the emphasis on human-directed use rather than autonomous contribution.

Drawing on the CIH, the interpretation of these examples can be analysed in terms of a single combined input in which source- and target-related structures are co-activated and organised. The process begins with cueing, whereby *AI* activates the target domain of computational systems associated with algorithmic processing, task execution, and data-driven computation. Through metonymic reduction, this domain is profiled in terms of its functional operation, foregrounding AI's capacity to perform tasks while backgrounding aspects such as intentionality, autonomy, or independent decision-making. The lexical item *tool* in Example (1) and the verb *use* in Example (2) cue the same source domain of instrumentality. Although realised through different lexical forms, both expressions activate a conceptual frame in which an artefact is employed by a user to accomplish a particular purpose. Metonymic reduction then profiles those aspects of the source domain most relevant to the context, foregrounding user control, instrumentality, and means-end relations while backgrounding properties unrelated to this

[1] Page numbers are included only for references to printed texts.

construal. The resulting combined input therefore brings together a functionally profiled representation of AI with an instrumentally profiled representation of a tool, providing the conceptual basis for the metaphorical construal of AI as an object employed to achieve human goals.

At a higher level of abstraction, the metonymically profiled domains converge on a shared means-end organisation, whereby actions are construed as serving externally defined goals. This shared organisation licenses the correlation between the foregrounded elements of the AI and tool domains, establishing a conceptual correspondence grounded in functional dependency rather than cognition or autonomous agency.

Within the combined input, these profiled structures are integrated into a coherent construal of AI as an instrument employed by a human user to achieve specific purposes. Metaphorical projection therefore yields an interpretation in which AI is characterised by controllability, predictability, and limited autonomy. This configuration, in turn, supports affective inferences associated with reassurance, low perceived risk, and instrumental trust, while post-propositional implicatures, such as reliability, safety, and the attribution of responsibility to the user, arise through inference over the integrated structure. In this construal, metonymic profiling plays a decisive role in downscaling AI by restricting the range of properties available for metaphorical projection. Figure 1 summarises this process.

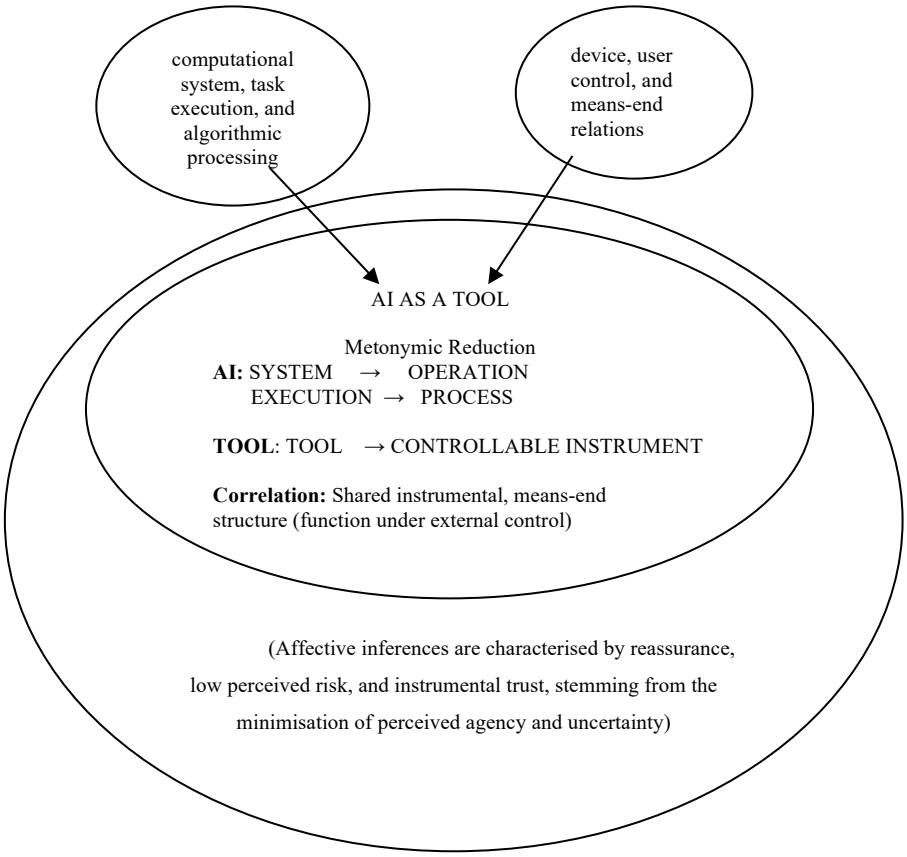

Figure 1. The CIH analytical process illustrated through the conceptual metaphor AI IS A TOOL

In this way, AI IS A TOOL occupies the lowest level of conceptual scaling because it attributes minimal agency, autonomy, and interpretive complexity to AI while maximising human control.

However, a shift in which elements of the tool-based construal are highlighted can scale up the representation by introducing limited agency, as demonstrated by (3).

(3) *Use AI to generate prompts or explore different writing styles to keep your creativity flowing.* (Sira, 2024)

In this example, AI continues to be construed as a tool employed by a human user. However, the interpretation shifts from task execution towards creative support. Through metonymic reprofiling, attention moves from execution-oriented functionality (TASK FOR OPERATION) to idea generation and exploratory assistance (OPERATION FOR CREATIVE SUPPORT), foregrounding the role of AI in shaping the creative process rather than merely carrying out predefined actions. Although human control is maintained through the imperative *use*, the verbs *generate* and *explore* profile activities associated with idea generation and creative exploration, contributing to a more agentive construal of AI than examples that foreground task execution alone. This incremental shift illustrates how conceptual scaling can occur within a single metaphor through changes in metonymic profiling, providing a basis for more elaborated role-based construals, such as AI IS A CO-PILOT, in which decision support and collaboration are further organised in terms of coordinated action under asymmetrical control.

5.2. AI IS A CO-PILOT

In contexts where AI is construed as providing specialised capabilities and broadening access to expert knowledge, as in (4) and (5), the system is positioned as a collaborative partner whose contribution extends beyond task execution to cognitive support and domain expertise.

(4) *… we can navigate this transformative era, ensuring that AI serves as a valuable co-pilot, enhancing human capabilities while respecting our agency and expertise.* (Sellen & Horvitz, 2024, p.23)

(5) *... co-pilots promise to democratize access to expert-level knowledge across professional domains.* (Perez, 2025)

Drawing on the Combined Input Hypothesis, the interpretation of these examples can be analysed as the progressive construction of a particular co-pilot construal. The process begins with cueing, whereby the lexical items *AI* and *co-pilot*, together with the surrounding discourse, activate their respective conceptual domains. In (4), expressions such as *enhancing human capabilities* and *respecting our agency and expertise* foreground AI as extending human competence while preserving human authority. In (5), *democratize access to expert-level knowledge* profiles AI as making specialised expertise more widely available. Through metonymic profiling, computational processes, large-scale knowledge processing, and predictive capabilities provide access to a reduced conceptualisation in which AI is construed primarily as a source of expert support and advanced capability (SYSTEM FOR SPECIALISED CAPABILITY; KNOWLEDGE PROCESSING FOR EXPERTISE), while autonomous intention and independent decision-making remain backgrounded.

The lexical item *co-pilot* simultaneously activates an aviation frame structured around complementary but asymmetrical roles. Through metonymic profiling, this frame is reduced to collaborative assistance under human authority (PILOT-CO-PILOT SYSTEM FOR HIERARCHICAL COLLABORATION; COORDINATION FOR EXPERT SUPPORT), while independent command is backgrounded. At a higher level of abstraction, the profiled elements of both domains converge on a shared relational structure in which one participant contributes specialised knowledge and capabilities that complement, rather than replace, the expertise of another. This shared structure sanctions the correlation underlying metaphorical projection, yielding a construal of AI as a

knowledgeable collaborator that broadens access to expertise while remaining accountable to human judgement.

Compared with the earlier AI IS A TOOL metaphor, the expressions in (4) and (5) scale AI upwards by attributing greater capability and operational agency, not because the system is represented as acting autonomously, but because it is construed as contributing expertise that actively informs human activity. This gives rise to affective and evaluative inferences such as calibrated trust, confidence, and conditional reliance, which emerge as post-propositional implicatures from the integrated conceptual structure.

Within the same co-pilot framing, however, further shifts in metonymic profiling can produce more elaborated construals. In (6), foregrounding AI's perceived contribution to human performance rather than specialised expertise alone produces a more highly scaled construal of the co-pilot metaphor.

(6) *As we stand on the threshold of this invisible co-pilot revolution, perhaps the most remarkable insight is how natural it feels. These systems don't make us feel like cyborgs or robots. They make us feel like the best possible versions of ourselves — confident, capable, and ready for whatever challenge comes next. And maybe that's the point.* (Perez, 2025)

In this construal, the interpretation shifts from assistance-oriented functions to augmentation-oriented functions (SUPPORT FOR ENHANCEMENT; TASK ASSISTANCE FOR ENHANCED PERFORMANCE). Through metonymic reprofiling, attention moves from the immediate tasks performed with AI to the broader effects attributed to its involvement, foregrounding confidence, capability, and readiness as outcomes associated with co-pilot interaction. The metaphor therefore scales AI upwards by representing the system as contributing to how users perform and experience activities rather than merely supporting the completion of individual tasks. At the same time, the description of AI as an *invisible* co-pilot foregrounds the seamless integration of the technology into human activity, suggesting that its contribution is experienced without disrupting the user's sense of agency, control, and confidence. This seamlessness supports evaluative inferences of calibrated trust in the system.

While the emphasis here is on the naturalness of the interaction rather than on technological opacity, the profiling of AI as operating in the background introduces a conceptual feature that becomes central in the subsequent AI IS A BLACK BOX construal, where hidden internal processes shape interpretations of the system's capabilities and limitations. The example therefore illustrates how conceptual scaling may involve not only changes in attributed capability and agency but also the redistribution of shared conceptual features across neighbouring metaphorical construals.

5.3. AI IS A BLACK BOX

The AI IS A BLACK BOX metaphor contrasts with the previous framings by shifting the basis of interpretation from capability and agency to epistemic accessibility. As shown by the examples (7) and (8), the metaphor directs attention away from what AI contributes towards the hidden processes through which those contributions are produced.

(7) *This lack of transparency is what makes AI a "black box" technology.* (Yadav, 2025)

(8) *The rise of generative AI, however, has embedded these systems into critical areas of daily life, making it essential to "open the hood" and examine what lies inside the black box.* (Yadav, 2025)

The process of interpreting these examples begins with the expressions *AI* and *black box*, together with contextual knowledge, activating relevant conceptual material. The term *AI* evokes computational systems characterised by algorithmic processing and input-output transformation.

Within this framing, metonymic reduction narrows attention to the inaccessibility of internal mechanisms (SYSTEM FOR INTERNAL PROCESS; OPERATION FOR HIDDEN PROCEDURE), while features such as performance, adaptability, or interaction recede from view. The expression *black box* in turn introduces a complementary structuring device grounded in the CONTAINER image schema, which organises experience in terms of an interior, a boundary, and an exterior (Lakoff & Johnson, 1999). The modifier *black* intensifies this construal by foregrounding opacity, restricted visibility, and limited access to what lies inside. Together, these elements support an interpretation in which inputs and outputs remain observable, but the processes linking them are concealed from view.

At a higher level of abstraction, the metonymically reduced AI domain and the container-based structure of the "black box" align through a shared representation of input-output processing with hidden internal states. This establishes correlation by mapping the inaccessibility of a sealed container's interior onto the limited interpretability of AI's computational operations. At the same time, contrast plays a delimiting role in this instance: unlike the co-pilot construal, which foregrounds AI's contribution to coordinated human activity, the black box metaphor shifts attention away from capability and interaction towards uncertainty about how outputs are produced.

Together, these examples yield a construal of AI as a complex system whose outputs are available for observation and use, but whose internal operations remain difficult to inspect or explain. Scaling therefore operates differently from the previous metaphors: rather than increasing or reducing attributed agency directly, the black box construal amplifies perceived complexity, uncertainty, and epistemic distance. This gives rise to affective and evaluative inferences, including uncertainty, reduced trust, perceived risk, and concerns about accountability, which emerge as post-propositional implicatures from the integrated conceptual structure. As (9) illustrates, further shifts in metonymic profiling can elaborate this epistemic construal by foregrounding the broader consequences of opacity.

(9) *This "black box" nature of AI means decisions and outputs are often opaque, even to experts, making effective regulation difficult.* (Yaldav, 2025)

In this example, metonymic profiling shifts from hidden computational processes (SYSTEM FOR CONCEALED MECHANISM) to the practical implications of that concealment (HIDDEN MECHANISM FOR LIMITED INTERPRETABILITY; OPAQUE PROCESS FOR REGULATORY CHALLENGE). Rather than introducing new conceptual content, the example reorganises the existing black box construal by foregrounding the widening gap between AI's observable effects and human capacity to understand, evaluate, or regulate them. This produces a more elaborated construal in which epistemic opacity is no longer merely a characteristic of AI systems but a factor shaping institutional oversight, accountability, and public trust.

The progressive amplification of epistemic distance in the black box construal also provides a basis for more highly scaled metaphorical construals, such as AI IS AN ALIEN INTELLIGENCE, in which opacity is no longer profiled simply as hidden computation but as evidence of a fundamentally different mode of cognition, further scaling AI's perceived capabilities and interpretive distance from human intelligence.

5.4. AI IS AN ALIEN INTELLIGENCE

Unlike AI IS A BLACK BOX, where interpretation is organised around epistemic inaccessibility and the suspension of interpretive access to internal processes, AI AS AN ALIEN INTELLIGENCE extends the scaling trajectory by shifting attention from hidden computation to radically non-human cognition. Rather than construing opacity as a limitation, the metaphor reinterprets it as evidence of forms of intelligence that exceed familiar human cognitive frameworks, resulting in a pronounced upward scaling of AI representations, as illustrated in (10).

(10) *... a provocative theory suggests these advanced systems might spawn alien artificial intelligence that breaks all known rules.* (Kesiena, 2025)

The interpretation of this example begins with the expression *AI* activating relevant conceptual material from the target domain, including computational systems, machine learning, and artificial forms of information processing, while the metaphorical cue *alien* activates a source domain associated with radical otherness and unfamiliarity. Through metonymic profiling, attention within the AI domain is narrowed from the broader technological system to capacities for generating forms of cognition that extend beyond familiar human modes of reasoning (SYSTEM FOR COGNITIVE CAPABILITY; COMPUTATIONAL PROCESS FOR NON-HUMAN INTELLIGENCE), while application-specific functionality and routine task execution are backgrounded. The adjective *alien* simultaneously profiles radical cognitive otherness rather than merely physical foreignness, foregrounding difference from established human patterns of thought. The expression *breaks all known rules* further elaborates this construal by highlighting a departure from existing conceptual frameworks, suggesting that familiar assumptions about intelligence are no longer sufficient for interpreting AI.

At a higher level of abstraction, the metonymically profiled domains converge on a shared conceptual structure of intelligence operating beyond familiar human parameters. This establishes correlation not through shared roles or functional coordination, as in the AI IS A CO-PILOT metaphor, nor through structural opacity, as in AI IS A BLACK BOX, but through the mapping of computational intelligence onto radical cognitive otherness. Whereas the "black box" metaphor construes limited interpretability as an epistemic boundary, the alien intelligence construal transforms that boundary into evidence of qualitatively different forms of cognition.

The resulting construal presents AI as a fundamentally non-human intelligence whose reasoning processes, capabilities, and potential modes of operation extend beyond established human understanding. This represents a high level of conceptual scaling, in which perceived capability, autonomy, and cognitive scope are amplified simultaneously. Affective and evaluative inferences, including awe, unease, fascination, and existential concern, emerge as post-propositional implicatures derived from this integrated conceptual structure.

However, the AI IS AN ALIEN INTELLIGENCE metaphor remains sensitive to shifts in metonymic profiling. Whereas (10) foregrounds AI's radical cognitive otherness by presenting it as an intelligence that *breaks all known rules*, (11) foregrounds the practical value of that otherness by shifting attention from the unfamiliarity of AI's internal architecture to the superiority of its outcomes.

(11) *... the specific mechanisms producing intelligence matter less than the outcomes achieved. If alien intelligence delivers superior results, the unfamiliar architecture becomes irrelevant.* (Kesiena, 2025)

In this example, metonymic profiling shifts from non-human cognitive architecture (SYSTEM FOR ALIEN COGNITION) to its observable consequences (ALIEN COGNITION FOR SUPERIOR PERFORMANCE). Rather than treating unfamiliarity as an obstacle to understanding, the example backgrounds concerns about interpretability and foregrounds the effectiveness of AI's outputs. The resulting construal therefore further scales up AI by presenting its cognitive difference as a source of enhanced capability rather than merely radical otherness. This demonstrates that the AI IS AN ALIEN INTELLIGENCE metaphor is not fixed but varies according to which aspects of the computational domain are made salient in context, producing different evaluations of the same non-human intelligence.

This progressive scaling also provides a basis for further metaphorical developments. When AI is construed as a form of intelligence that exceeds familiar human cognitive frameworks, shifts in metonymic profiling can redirect attention from exceptional capability to its broader implications, including concerns about human control, alignment, vulnerability, and ultimately existential risk.

5.5. AI IS AN EXISTENTIAL RISK

The AI IS AN EXISTENTIAL RISK metaphor represents the upper extreme of the proposed scaling continuum because it involves a shift from representing AI in terms of its properties or roles to construing its development through metaphors of escalation, uncontrollability, and existential threat, as illustrated in (12).

(12) *Creating intelligence that operates on completely unknown principles could amplify these risks exponentially, introducing dangers that humans cannot anticipate or counter.* (Kesiena, 2025)

The interpretation of this example begins with cueing, whereby references to creating intelligence activate relevant conceptual material from the AI domain, while expressions referring to amplified risks and human inability to anticipate or counter consequences activate the risk domain. The reference to *creating intelligence* cues advanced AI systems characterised by increasing computational complexity and cognitive performance. Through metonymic profiling, attention is directed towards the emergence of forms of intelligence operating beyond familiar human frameworks (SYSTEM FOR ADVANCED COGNITION; COMPUTATIONAL PROCESS FOR UNFAMILIAR OPERATIONS), while application-specific functionality and bounded task performance are backgrounded. The expression *completely unknown principles* further reinforces this construal by foregrounding modes of operation that cannot be readily interpreted through existing human knowledge.

The second part of the example redirects attention from the characteristics of the intelligence itself to the consequences associated with its operation. Expressions such as *amplify these risks exponentially* and *humans cannot anticipate or counter* metonymically foreground escalation, unpredictability, and limitations of human intervention (RISK FOR AMPLIFICATION DYNAMICS; CONSEQUENCE FOR CASCADING EFFECTS; HUMAN RESPONSE FOR INABILITY TO PREDICT OR MANAGE), while more localised or reversible outcomes recede from view.

At a higher level of abstraction, the metonymically profiled domains converge around a metaphorical structure in which unfamiliar AI capabilities are conceptualised through the logic of escalation and loss of containment. This establishes correlation between computational processes operating beyond established human understanding and scenarios of systemic vulnerability. The resulting construal presents AI not primarily as an agent with greater autonomy, but as a form of intelligence whose unfamiliar operations are associated with consequences that may exceed human capacities for prediction and intervention. Scaling therefore occurs through the expansion of projected consequences rather than through the attribution of additional agency. Affective and evaluative inferences, including urgency, fear, vulnerability, and existential concern, emerge as post-propositional implicatures derived from this integrated conceptual structure.

The existential-risk construal can be further elaborated through changes in metonymic profiling, as illustrated in (13).

(13) *The possibility of multiple intelligence types existing simultaneously raises additional questions about how they might interact, complement or conflict with each other in shaping humanity's future.* (Kesiena, 2025)

Whereas (12) foregrounds the consequences of a single form of unfamiliar intelligence, (13) shifts attention to the interactions among multiple intelligences. Through metonymic profiling, emphasis moves from individual capability to the relationships between different forms of intelligence (INTELLIGENCE FOR PLURALITY; PLURALITY FOR INTERACTION). The resulting construal expands the scope of interpretation from the consequences of a single system to the complexity of multiple interacting intelligences and their implications for humanity's future.

Rather than further increasing attributed capability, the metaphor scales upwards by broadening the systemic scope and long-term consequences of the projected scenario.

This upper limit of conceptual scaling is subsequently challenged by metaphorical construals that question whether AI possesses the forms of intelligence attributed to it. The AI IS A STOCHASTIC PARROT metaphor represents such a counter-scaling by reinterpreting apparently intelligent behaviour as the product of statistical pattern reproduction rather than genuine understanding.

5.6. AI IS A STOCHASTIC PARROT

Construing AI with phrases such as *stochastic parrot* contrasts with construals such as co-pilot and alien intelligence by challenging attributions of understanding, intentionality, and cognitive depth, resulting in a strong downward scaling of AI representations, as (14) demonstrates.

(14) *... an LM is a system for haphazardly stitching together sequences of linguistic forms it has observed in its vast training data, according to probabilistic information about how they combine, but without any reference to meaning: a stochastic parrot.* (Bender et al., 2021, p. 617)

The interpretation of this example begins with cueing, whereby expressions associated with language modelling, linguistic forms, and probabilistic generation activate relevant conceptual material from the AI domain, while the metaphorical expression *stochastic parrot* introduces a source-domain structure associated with imitation and reproduced speech.

Through metonymic profiling, attention is narrowed from the broader system to its statistical operations (SYSTEM FOR PROBABILITY-BASED GENERATION; LINGUISTIC PRODUCTION FOR FORM PREDICTION), while capacities associated with meaning, understanding, and intentional communication are backgrounded. The expression *stochastic* further foregrounds probabilistic combination rather than deliberate selection, reinforcing the interpretation of AI outputs as the result of statistical processes rather than meaningful production.

The source domain of *parrot* activates a frame associated with vocal imitation and reproduced expression. Through metonymic reduction, this domain is narrowed to surface-level repetition (ANIMAL FOR IMITATIVE BEHAVIOUR; VOCALISATION FOR REPRODUCED FORMS), while features such as comprehension, intention, or communicative purpose are backgrounded. The mapping of these selectively profiled domains supports a construal of AI as producing linguistic outputs through imitation-like processes without semantic understanding.

At a higher level of abstraction, the reduced domains converge around a shared structure of form production without meaningful comprehension. This establishes correlation between probabilistic language generation in the AI domain and imitative reproduction in the parrot domain, yielding the interpretation of AI as a system capable of producing convincing linguistic patterns without possessing the understanding associated with human communication. In contrast to upward-scaled metaphors that foreground expertise, cognition, or exceptional capability, the stochastic parrot construal redirects attention to the limits of AI representation by separating successful output generation from underlying understanding. This configuration gives rise to affective and evaluative inferences such as scepticism, reduced trust, and demystification, which emerge as post-propositional implicatures from the integrated conceptual structure.

Importantly, the AI IS A STOCHASTIC PARROT metaphor is not associated with a single fixed construal. Changes in metonymic profiling can selectively foreground different aspects of the source domain, yielding alternative interpretations while preserving the same underlying metaphor. Consider the use of the metaphor in (15).

(15) *By treating the AI not as a servant but as a cohabitant with character, the Stochastic Parrot opens the door to richer, more spontaneous, and more emotionally textured ways of living with intelligent machines. (*Chang, Kuang & Coelho, 2025).

Whereas Example (14) metonymically profiles the parrot primarily in terms of imitation and reproduced speech, (15) foregrounds different aspects of the same source domain by construing the parrot as a cohabitant with character. This shift expands the source-domain structure from imitative behaviour to companionship and social presence (PARROT FOR COMPANION; COMPANION FOR COHABITATION; COHABITATION FOR EMOTIONALLY TEXTURED INTERACTION). Through this metonymic reconfiguration, the metaphor no longer serves primarily to question AI's understanding but to foreground its capacity to participate in meaningful human relationships. The resulting construal therefore scales up the relational and affective dimensions of the metaphor while leaving its underlying cognitive limitations intact.

5.7. Summary

The CIH analysis presented in this section demonstrates that metaphorical construals of AI are not fixed. Changes in metonymic profiling reorganise the conceptual material available for metaphorical projection, allowing the same metaphor to support different interpretations depending on the discourse context. As illustrated by the variation within the AI IS A STOCHASTIC PARROT metaphor, such reconfiguration may involve not only changes in perceived capability or understanding, but also shifts in the relational and affective dimensions through which AI is conceptualised.

A broader implication of this analysis is that the evolution of figurative discourse about AI can be understood as an ongoing process of conceptual scaling, whereby different metaphorical construals selectively amplify, attenuate, or reorganise particular dimensions of AI, including capability, agency, epistemic status, and relational significance. Emerging metaphors do not simply replace existing ones; rather, they reflect alternative selections and organisations of conceptual material that make different aspects of the technology salient. For example, the recent framing *AI's Not a Genie in a Lamp: It's a Space to Think* (The Presidents' Forum, 2025) illustrates how alternative metaphorical construals can recalibrate understandings of AI. The metaphor AI IS A GENIE IN A LAMP profiles AI's capacity to generate solutions with apparent ease and projects these features onto a source domain associated with magical power, instant fulfilment, and extraordinary capability. Through this projection, AI is construed as possessing enhanced problem-solving potential and a degree of exceptional capability. By contrast, AI IS A SPACE TO THINK foregrounds AI as an enabling environment for reflection, exploration, and human creativity, projecting features of openness and support rather than autonomous power. This construal maintains human agency while repositioning AI as a resource for extending human activity.

This contrast demonstrates the value of the CIH as a framework for analysing the ongoing development of AI metaphors. It has the potential to show how changing patterns of metonymic profiling and metaphorical projection generate alternative conceptualisations of the technology, not through the replacement of one metaphor by another, but through the selective reorganisation of conceptual material that shapes how AI capabilities, roles, and relationships with humans are understood.

6. Conclusion

The findings reported in this article demonstrate that AI discourse is shaped not by isolated metaphors or metonymies, but by hybrid figurative constructions in which metaphor and metonymy interact to produce scale-sensitive conceptualisations of AI. By applying the Combined Input Hypothesis, it shows that metonymic operations function as primary mechanisms for regulating interpretive scope. These operations determine which aspects of the AI domain are made salient and thus what counts as the referent in a given context, while metaphorical mappings

organise this selected material into structured configurations that support inferences about agency, capability, and impact.

A key contribution of the study is demonstrating that variation in representations of AI arises not only across different conceptual metaphors, but also within individual metaphorical framings. The analysis shows that the same metaphor can support divergent construals depending on how conceptual material is selectively activated and organised in context. This finding shifts the focus from metaphor as a fixed mapping between domains to metaphor as a dynamic process in which meaning emerges through the interaction of metonymic profiling and metaphorical projection. More broadly, the study reconceptualises scaling in discourse about AI as an inferential effect produced through these interacting processes. Rather than reflecting rhetorical exaggeration or inherent properties of particular metaphors, differences in how AI is construed in terms of capability, relation to human agency, epistemic accessibility, and potential consequences arise from systematic variations in which aspects of the technology are foregrounded, backgrounded, and integrated into broader interpretive structures. In this way, the study provides a unified account, grounded in the CIH, of how figurative language contributes to both conceptualising AI and generating the evaluative and affective responses associated with it.

The continuing evolution of AI discourse further demonstrates the value of this account. Emerging framings do not simply introduce new representations of the technology. Rather, they reflect changing patterns of metonymic profiling and metaphorical projection through which the perceived capabilities and role of AI are continually recalibrated. The CIH therefore provides a principled framework not only for analysing existing metaphorical construals, but also for explaining how future AI metaphors may emerge as public discourse negotiates changing understandings of the technology.

A limitation of the present study is that it is based on a purposive sample of examples representing scaling in AI discourse. This sampling strategy was appropriate for the study objective of providing a systematic qualitative account of the role of metaphor and metonymy in scaling. However, it does not permit conclusions about the prevalence or distribution of the identified patterns across AI discourse more generally. These questions would be more appropriately addressed through a complementary corpus-based investigation.

Acknowledgements

I am grateful to the anonymous reviewer and the PJOS editors for their thoughtful and constructive feedback that has helped strengthen this manuscript. Any remaining weaknesses are my responsibility.

References

Barcelona, A. (2000). On the plausibility of claiming a metonymic motivation for conceptual metaphor. In A. Barcelona (Ed.). *Metaphor and metonymy at the crossroads: A cognitive perspective* (pp. 31-58). Berlin: Walter de Gruyter Mouton.

Barcelona, A. (2015). Metonymy. In E. Dabrowska and D. Divjak (Eds.). *Handbook of cognitive linguistics* (Vol. 39) (pp. 143-167). Berlin: De Gruyter Mouton.

Bender, E. M., Gebru, T., McMillan-Major, A. and Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? FaccT'21: Proceedings of the *2021 ACM Conference, Virtual Event Canada*, 3-10 March 2021 (pp. 610-623). New York: ACM. https://doi.org/10.1145/3442188.3445922

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

Campolo, A. and Crawford, K. (2020). Enchanted determinism: Power without responsibility in artificial intelligence. *Engaging Science, Technology, and Society* 6: 1-19. https://doi.org/10.17351/ests2020.277

Cave, S. and 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

Chang, E., Kuang, Q. and Coelho, M. (2025). *The stochastic parrot: A physical AI cohabitant*. MIT Design Intelligence. https://designintelligence.mit.edu/work/stochastic-parrot

Cheng, M., Lee, A. Y., Rapuano, K., Niederhoffer, K., Liebscher, A. and Hancock, J. (2025). From tools to thieves: Measuring and understanding public perceptions of AI through crowdsourced metaphors. *arXiv.* https://doi.org/10.48550/arXiv.2501.18045

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

Croft, W. (2008). On iconicity of distance. *Cognitive Linguistics* 19(1): 49-57. https://doi.org/10.1515/COG.2008.003

Emard Sira, M. (2024, March 5). *How to use artificial intelligence in a sentence? Easy examples*. Yorkville College. https://www.yorkvillecollege.com/artificial-intelligence-in-a-sentence

Fauconnier, G. and Turner , M. (2002). *The way we think*. New York: Basic Books.

Goossens, L. (1990). Metaphtonymy: The interaction of metaphor and metonymy in expressions for linguistic action. *Cognitive Linguistics* 1(3): 323-340. https://api.semanticscholar.org/CorpusID:144359761

Harari, Y. N. (2023). AI and the future of humanity. *Frontiers Forum Live*, 31 de mayo. Disponible en: https://www.youtube.com/watch?v=azwt2pxn3UI

Inie, N., Zukerman, P., & Bender, E. M. (2026). *De-anthropomorphizing "AI": From wishful mnemonics to accurate nomenclature*. First Monday. https://firstmonday.org/ojs/index.php/fm/article/view/14366/12165

Johnson, D.G. and Verdicchio, M. (2017). Reframing AI discourse. *Minds & Machines* 27: 575-590. https://doi.org/10.1007/s11023-017-9417-6

Kesiena, D. (2025, October 18). *The terrifying truth about alien AI taking over*. Rolling Out. https://rollingout.com/2025/10/18/the-truth-about-alien-ai-taking-over/

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

Lakoff, G. and Johnson, M. (1999). *Philosophy in the flesh: The embodied mind and its challenge to Western thought*. New York: Basic Books.

Langacker, R. W. (1986). An introduction to cognitive grammar. *Cognitive Science* 10(1): 1-40. 10.1016/S0364-0213(86)80007-6

Lawrence, S. (2025). *Stop treating AI like a tool*. Sadie's Newsletter. https://sadiestlawrence.substack.com/p/stop-treating-ai-like-a-tool

Maas, M. M. (2023). AI is like… A literature review of AI metaphors and why they matter for policy. *SSRN Electronic Journal*. http://dx.doi.org/10.2139/ssrn.4612468

Matthews, S. W. and Danesi, M. (2019). AI: A semiotic perspective. *Chinese Semiotic Studies* 15(2): 199-216. https://doi.org/10.1515/css-2019-0013

Mishra, P. and 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

Natale, S. and Ballatore, A. (2020). Imagining the thinking machine: Technological myths and the rise of artificial intelligence. *Convergence* 26(1): 3-18. https://doi.org/10.1177/1354856517715164

Nerlich, B. (2025a). *Metaphors for AI: An overview of recent studies*. Making Science Public. https://makingsciencepublic.com/2025/11/21/ai-metaphor-studies-an-overview/

Nerlich, B. (2025b). *Metaphors for AI: What changed and why it matters*. Making Science Public. https://makingsciencepublic.com/2025/11/28/observing-shifts-in-metaphors-for-ai-what-changed-and-why-it-matters/

Nerlich, B. (2026). *Metaphors in AI and genomics: Going beyond blueprints and parrots*. Making Science Public. https://makingsciencepublic.com/2026/03/13/

Oster, N., McCaleb, L. and Mishra, P. (2025). Swiss army knives, stochastic parrots, drunk interns, and overlords: Understanding AI through metaphors. In *Society for Information Technology & Teacher Education International Conference* (pp. 842-845). Association for the Advancement of Computing in Education (AACE).

Perez, C. E. (2025, June 22). *The invisible co-pilot revolution: How real-time AI is redefining human performance. Intuition Machine.* Medium. https://medium.com/intuitionmachine/the-invisible-co-pilot-revolution-how-real-time-ai-is-redefining-human-performance-a16d66527468

Ruiz de Mendoza Ibáñez, F. J. R. (1998). On the nature of blending as a cognitive phenomenon. *Journal of Pragmatics* 30(3): 259-274.

Ruiz de Mendoza Ibañez, F. J. and Peña Cervel, S. (2002). Cognitive operations and projection spaces. Jezikoslovlje 3(1-2): 131-158. http://hrcak.srce.hr/file/49697

Ruiz de Mendoza Ibañez, F. J. and Diéz, O. I. (2002). Patterns of conceptual interaction. In R. Pörings and R. Dirven (Eds.). *Metaphor and metonymy in comparison and contrast* (pp. 501-46). Berlin: Mouton de Gruyter.

Ruiz de Mendoza Ibañez, F. J. and Pérez, L. (2003). Cognitive operations and pragmatic implication. In K. Panther and L. Thornburg (Eds.). *Metonymy and pragmatic inferencing* (pp. 23-49). Amsterdam: John Benjamins.

Ruiz de Mendoza Ibáñez, F. J. R. and Sáenz, F. S. (2003). Content and formal cognitive operations in construing meaning. *Italian Journal of Linguistics* 15: 293-320.

Sellen, A. and Horvitz, E. (2024). *The rise of the AI co-pilot: Lessons for design from aviation and beyond*. *Communications of the ACM* 67(7): 18-23. https://doi.org/10.1145/3637865

Sperber, D. and Wilson, D. (1995). *Relevance: Communication and cognition* (2nd edn.). Oxford: Blackwell.

The Presidents' Forum. (2025, February 6). *AI's not a genie in a lamp: It's a space to think*. President Forum. https://presidentsforum.org/2025/02/06/ais-not-a-genie-in-a-lamp-its-a-space-to-think/

Vallis, C., Wilson, S. and Casey, A. (2025). Fear and awe: Making sense of Generative AI through metaphor. *Journal of Interactive Media in Education* 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: Oxford University Press.

Woods, K. (2025). If AI is our co-pilot, who is the captain? *AI & Society* 40(3): 1537-1538. https://doi.org/10.1007/s00146-024-01965-5

Yadav, V. (2025, August 18). *AI black box: Problem, challenges, consequences & regulation solutions*. VIESTORIES. https://viestories.com/insight/expert-corner/ai-black-box-problem-challenges-consequences-regulation-solutions-9672704

Zirenko, M., Machura, I. A., Fabriz, S., Schulze-Vorberg, L. and Horz, H. (2025). AI and learning with AI: University students' metaphorical conceptualizations. *Journal of Interactive Media in Education* 1. https://doi.org/10.5334/jime.994.

Zlatev, J. and Moskaluk, K. (2022). Translation validity in metaphor theories: CMT, DMT and the Motivation & Sedimentation Model. In A. Bagasheva and B. Hristov (Eds.). *Figurativity and human ecology* (pp. 123-149). Amsterdam: John Benjamins. https://doi.org/10.1075/ftl.17

Ghsoon Reda is a postgraduate research supervisor at the University of Wales Trinity Saint David. She earned a PhD in Applied Linguistics from Leicester University and an MA in Theoretical Linguistics from Sussex University, UK. Her research interests include Semiotics, Cognitive Semantics, Pragmatics, Educational Psychology, and English Language Teaching.