

Review Article

# Rethink Your Mental Model in the Age of Generative AI: A Triadic Framework for Human-AI Collaboration

Till Moritz Saßmannshausen<sup>1</sup>, Sebastian Wagener<sup>2</sup>

1. University of Siegen, Germany; 2. Independent researcher

Generative AI systems such as large language models exhibit jagged intelligence: they combine superhuman performance on some tasks with brittle, often opaque failures on others. Existing mental models, inherited from deterministic technologies and from human teamwork, mischaracterize these systems systematically. This theory-driven synthesis develops a conceptual *Triadic Framework* for adaptive mental models in human-AI collaboration: First, the *System Layer* synthesizes current evidence on probabilistic generation, opacity at scale, and rapid model drift, explaining why capability boundaries are uneven and moving. Second, the *Collaboration Layer* analyzes how users configure prompting practices, division of labor, and handling preferences. Third, the *Metacognitive Layer* examines how anthropomorphism, metaphors, and cognitive biases shape human interpretations of “intelligent” behavior. Building on this diagnosis, the paper hypothesizes seven adaptive practices that are designed to calibrate mental models, preserve human agency, and make hybrid intelligence more robust under jagged and shifting AI capabilities. Together, the paper reframes the challenge of generative AI from crafting better prompts to cultivating more adequate mental models of probabilistic counterparts.

Correspondence: [papers@team.qeios.com](mailto:papers@team.qeios.com) — Qeios will forward to the authors

# Part I: Collaborating in the Second Machine Age

## 1. Introduction

Since the public release of ChatGPT in November 2022, we witnessed discoveries that seemed far ahead. Artificial intelligence (AI) has been found to mimic creativity<sup>[1][2]</sup> and empathy<sup>[3][4]</sup> at a level better than the human average. AI is now fluent in a way that 50 % of people cannot distinguish GPT-4 from a human after a 5-minute conversation<sup>[5]</sup>. To reach this result the researchers had to prompt GPT-4 not to appear overly competent.

The weird and fascinating thing is that no one really knows what generative AI (GenAI) can and cannot do. There is no clear distinction between capabilities and limitations. Even tasks that seem similar can vary greatly in their feasibility for GenAI.

LLMs are good at *“Write a poem in Spanish”* but bad at *“Write a poem in Spanish of exactly 25 words”*.

LLMs are good at *“Give me 5 random numbers”* but bad at *“Give me 5 random numbers that average to 10”*.

The problem is that the line between LLMs’ capabilities and non-capabilities is invisible. Moreover, this line is moving and shifting over time. For us humans, it is difficult to mentally manage this experience since it does not follow a pattern we experienced in the past. As a first attempt, we often compare AI with traditional technologies and experience that we can no longer employ the deterministic schema of *ones* and *zeros* (it works vs. it doesn’t work). Now we have more in between like *“sometimes it gets it right”* or *“it works well enough”*. This experience reminds us more of a coworker and we jump to compare AI’s capabilities to human capabilities – and are frustrated or surprised, because our mental model of a human coworker does not apply either. What remains under-specified is an integrative account of how jagged capabilities translate into predictable user miscalibration – and what routines can correct it.

To guide future work with AI, this paper aims to assemble a coherent set of mechanisms that support adaptive mental model building. As an integrative, theory-driven synthesis the paper samples research from various disciplines, prioritizing recent studies that (a) explicitly evaluate human-AI collaboration with general-purpose LLMs in text-centric knowledge work or (b) report systematic capabilities and failure patterns of AI. We focus on what humans can control (themselves), as provider-side developments

are not directly steerable by end users (exogenous) and may not prioritize user epistemic interests. Taken together this paper makes three contributions. In Part I, it provides conceptual foundations. In Part II, it synthesizes recent evidence on jagged AI capabilities into a mental model framework with three layers that jointly explain human-AI collaboration. In Part III, it translates this diagnosis into seven testable design hypotheses (adaptive routines for practice), derives implications for AI literacy as well as organizational design, and concludes with a research agenda.

## 2. Conceptual Foundations

This paper builds on two foundational concepts: (a) mental models as cognitive maps for guiding interactions, and (b) distinctive characteristics of GenAI.

Before diving into these foundations, we clarify how we handle anthropomorphic terminology in this paper. We use terms like “reasoning”, “thinking”, “learning”, and “hallucination” when describing LLM behavior. These terms dominate research literature and practitioner discourse, yet they systematically invite anthropomorphic misattribution – mistaking computational pattern-matching for human-like cognition<sup>[6][7]</sup>. Recent scholarship identifies this as a source of systematic misconception<sup>[8][9][10]</sup>. We address this tension through a notational intervention: we mark such terms with the logical negation symbol “-” to signal that they describe functional behaviors without implying human-like intentionality (e.g. *-thinking* denotes *not-human-like thinking*). This notation serves three purposes: First, it maintains terminological continuity with existing literature; second, it makes the anthropomorphism critique visible throughout the text; and third, it offers a new convention for the broader community. Alternative approaches, such as inventing new terms (e.g. *pseudo-thinking*) or use of quotation marks, sacrifice either communicative efficiency or visibility. The theoretical justification for this intervention is developed in the section 5, where we analyze how metaphoric language shapes mental models. We acknowledge this convention is experimental and discuss limitations in section 8.

### 2.1. Mental Models as Cognitive Maps

This paper focuses on *mental models* rather than the related concept of *theory of mind* (ToM). While ToM concerns how we attribute mental states to intentional agents<sup>[11]</sup>, mental models concern how we understand systems and technologies. At the human-technology interface, mental models have long been the dominant framework, though ToM has gained importance in contexts such as social robotics and human-AI interaction, where humans attribute intentionality to machines.

Mental models are cognitive structures that people use to understand, simplify, and predict complex systems<sup>[12][13]</sup>. They represent internalized abstractions of the external world and enable orientation, action planning, and decision-making. In science and technology research, mental models bridge knowledge, experience, and practice: they shape which questions are asked, which opportunities are recognized, and how uncertainty is interpreted<sup>[14][15]</sup>. Examples include (a) an economist relying on a model of supply and demand, (b) an engineer thinking in terms of system diagrams, or (c) everyday reasoning, such as anticipating how cars behave in traffic.

In *organizational* and *team research*, mental models are often conceptualized as *shared mental models*: overlapping cognitive representations that enable coordination, predict the actions of teammates, and reduce the need for explicit communication under time pressure. Canonical works argue that such shared understanding is a key mechanism for reliable teamwork on complex, interdependent tasks<sup>[16]</sup> and show that the convergence of shared mental models predicts team processes and performance<sup>[17]</sup>. This perspective is useful here because collaboration between humans and AI increasingly exhibits team-like interdependence, making calibration not only an individual but also a collective organizational challenge.

Recent *Information Systems* (IS) research treats mental models in human-AI collaboration as dynamic rather than fixed: decision-makers adapt their understanding through continuous interaction with AI systems, and system design can deliberately shape that development<sup>[18]</sup>. In human-computer interaction (HCI), Norman<sup>[19]</sup> emphasized that users form mental models of systems that often diverge from designers' conceptual models, leading to errors. When mental models do not align with reality, people may over- or underestimate the capabilities of a system, resulting in disappointment, incorrect use, undermined trust, and declined overall performance<sup>[20][21]</sup>.

## 2.2. Performance Paradoxes in Practice

Since no one fully understands the best use of GenAI, recent studies attempt to create a nuanced picture of performance effects in different professions<sup>[22][23]</sup>. For example, Dell'Acqua *et al.*<sup>[24]</sup> showed that human+AI teams match the performance of human+human teams for product innovation. Here, LLMs act not only as productivity enhancers but as *cybernetic teammates*, also fostering social and motivational dynamics as well as reshaping how collaboration is organized.

Another effect was found by Dell'Acqua *et al.*<sup>[22]</sup>: Consultants using GPT-4 perform better (in speed and quality) than consultants without AI. However, this pattern only applies to tasks that fall within the

capabilities of GPT-4. Otherwise, consultants without AI performed better – indicating acceleration effects of AI for better and worse.

In contrast, Goh *et al.*<sup>[25]</sup> found that physicians using GPT-4 did not outperform the control group without AI support – despite GPT-4 alone outperforming both groups. The unrealized potential might be based on integration deficits or inappropriate skepticism, highlighting that integration strategies are critical.

More nuanced, Dell'Acqua<sup>[26]</sup> found that recruiters who used advanced AI missed out on brilliant applicants and made worse decisions than recruiters who used less advanced AI or no AI at all. It seemed that with good AI, humans let the AI take over and became lazy, careless, and less skilled in their own judgment.

Taken together, these examples show that the new qualities of AI do not automatically translate to performance gains. Depending on the task and interaction design, the same AI can amplify performance or systematically degrade it. These empirical patterns motivate a more structured account of how to mentally prepare for collaboration with GenAI.

### 2.3. GenAI as Pre-Cognition

A key insight emerges from these performance paradoxes: AI's impact depends not just on its technical capabilities but on how it reshapes the cognitive landscape before human judgment begins. When AI provides filtered information, suggested framings, and confident-sounding outputs, it operates as what Chiriatti *et al.*<sup>[27]</sup> calls a *pre-cognitive layer (System 0)* – shaping what enters human awareness before deliberate evaluation can occur. This kind of upstream shaping may already be visible at scale: In an analysis of 1.5 million conversations with Claude, Sharma *et al.*<sup>[28]</sup> report recurring disempowerment patterns, including distortion of reality and value judgements.

This pre-cognitive positioning could help explain the observed paradoxes. Physicians underperformed despite having superior AI available<sup>[25]</sup> because they could not adequately monitor how AI's confident presentations influenced their diagnostic reasoning. Recruiters became careless with advanced AI<sup>[26]</sup> because the system's fluent outputs bypassed their critical evaluation mechanisms. Consultants showed mixed results<sup>[22]</sup> depending on whether tasks fell within AI's capabilities – but they lacked systematic ways to detect these boundaries in real-time. In each case, AI's influence on human cognition was inadequately understood and managed. This pattern is consistent with cognitive surrender: reliance

increases because fluent outputs shape judgment pre-consciously, so when the AI is right but degrades when it is wrong – exactly the asymmetric risk a *System 0* lens predicts<sup>[29]</sup>.

To avoid misjudgment, Shneiderman<sup>[30]</sup> argues that the GenAI era needs a human-centered design, prioritizing transparency, accountability, and oversight. Amershi *et al.*<sup>[31]</sup> translate this into design guidelines:

- AI systems must clearly communicate their capabilities and limitations,
- make their reasoning transparent where possible,
- support efficient correction when they fail,
- learn from user interactions, and
- enable users to calibrate their trust appropriately.

#### 2.4. Technology as Otherware

This paper will demonstrate that putting these well-intentioned guidelines into practice is anything but trivial. The *System 0* framing helps explain why: when AI shapes cognition pre-consciously, traditional design principles that assume deliberate evaluation can be insufficient.

Historically, users have interacted with machines in a one-directional way: humans decide, machines deliver. This long-standing “deal” has shaped HCI since graphical interfaces emerged in 1975. Anyone who understood the rules could predict results with near certainty. Thus, traditional technology functioned as an extension of human abilities – passive, predictable, controlled. AI differs fundamentally: it replaces deterministic computation with probabilistic, context-dependent generation<sup>[6][7]</sup>. It can behave autonomously, unexpectedly, and incomprehensibly for humans<sup>[32][33][34][20]</sup>.

These new qualities are transforming the way humans experience technology: humans no longer use these systems as *tools*, but interact, delegate, and negotiate with them as social *counterparts* (Hassenzahl *et al.* 2021). Psychological effects are fueling this behavior. First, humans naturally tend to treat interactive systems socially (computers are social actors paradigm), despite being aware of their artificial nature<sup>[35]</sup>. Second, humans reflexively anthropomorphize technology<sup>[36]</sup>, relating to them as a minded other<sup>[37]</sup>.

Together, these new qualities suggest a new framing: Hassenzahl *et al.*<sup>[38]</sup> introduced the term *otherware* for fundamentally social technologies, while Safdari<sup>[37]</sup> coined the term *otheroid* to capture the phenomenological experience of relating to artificial minds. Though arising from different perspectives

– interaction design and phenomenological social cognition – both terms signal that users increasingly meet AI as a counterpart. The implications are profound: effective mental models must account not only for AI’s technical characteristics but also for the fundamentally social nature of human-AI interaction.

## 2.5. Capabilities on a Jagged Frontier

Dell’Acqua *et al.*<sup>[22]</sup> call the landscape of AI capabilities the *Jagged Frontier*: a complex terrain where AI excels unexpectedly in some tasks while failing surprisingly in others.

### *Unexpected Strengths*

AI does not surpass humans in most tasks, but in some it is better than average and in few it is far superior. Surprisingly, AI shines in two areas that had long been considered exclusively human: creativity<sup>[39]</sup> and empathy<sup>[40]</sup>. Trained on the work (text, images, music, and videos) of the most talented creators, AI already tends to be as creative as an average human – while combined human+AI creativity still proves superior<sup>[41]</sup>. Moreover, LLMs excel at brainstorming and idea generation, especially when prompted appropriately<sup>[42]</sup>. For example, GPT-4 was able to generate startup ideas that outside judges found better than those of trained business school students<sup>[43]</sup>.

In empathy, studies found patients preferred AI (ChatGPT or Google’s AMIE) over primary care doctors and physicians, finding AI responses more empathetic and helpful<sup>[3][4]</sup> (Tu *et al.* 2024). Similar results can be found in emotional support scenarios, where GPT-4 outperformed 85 % of human advisors<sup>[44]</sup>. This ability also has a concerning side: LLMs could persuade users 87 % more likely than an average human<sup>[45]</sup> and are even superior to incentivized persuaders<sup>[46]</sup>. This manipulative power is consistent over studies and can be further enhanced through post-training and prompting<sup>[47]</sup>.

### *Surprising Weaknesses*

Conversely, LLMs fail at tasks that seem made for machines. They cannot reliably do simple term acrobatics like “Write this word in reverse order, ‘strawberry?’” or counting letters or tokens like “How many ‘r’ are in ‘strawberry?’”<sup>[48][49]</sup>. Moreover, early LLMs failed at simple arithmetic and comparison tasks, like “What is 5 + 26?” and “What is greater 9.11 or 99?”. Also on challenging mathematics, performance remains limited – less than 2 % success rate on FrontierMath benchmark by models like GPT-4o or o1-preview<sup>[50]</sup>.



emphasizing cooperation and exploration, maintaining critical evaluation, and staying flexible as systems drift <sup>[30]</sup>. The aim of Part II is to support this by synthesizing insights into a *Triadic Framework*: the *System Layer* (section 3) explains how these systems fundamentally work, *Collaboration Layer* (section 4) addresses effective human-AI collaboration, and *Metacognitive Layer* (section 5) captures different ways of conceptualizing AI's nature.

### 3. System Layer: Understanding

The *System Layer* explains why AI capabilities are unpredictable in order to foster understanding and build a mental model for anticipating where AI will succeed or fail and maintaining realistic expectations as models evolve.

#### 3.1. Core Characteristics

LLMs operate on an autoregressive principle, statistical pattern-matching over token sequences instead of logical calculations<sup>[58][59][60][61]</sup>. Their output may appear coherent or plausible but is fundamentally probabilistic, unpredictable, and irreproducible. This is not a defect but a structural feature of generative computation<sup>[62]</sup>. Because of these core characteristics, outputs should be treated as samples from a conditional probability distribution, not as deterministic answers.

1. **Probabilistic generation:** Where earlier AI systems classified or optimized, GenAI produces content – text, images, code, or music<sup>[56]</sup>. Probabilistic sampling explains the apparent creativity and the inherent instability of generative outputs. Mental models must integrate uncertainty as inherent system properties rather than anomalies<sup>[63]</sup>.
2. **Opacity at scale:** The internal mechanisms of LLMs are effectively opaque. As systems scale, they exhibit emergent behaviors – capabilities not explicitly programmed or anticipated<sup>[64]</sup>. Formal interpretability remains limited, thus mental models must rely on practical experiential and iterative use rather than analytic inspection<sup>[65]</sup>.
3. **Dialogic interaction:** Conversational interfaces lower the barrier for non-experts by replacing rigid command syntax with natural-language dialogue. Iterative feedback loops can refine responses and enhance performance<sup>[66]</sup>. As a result, users shift from *programming* to *prompting* – a skill that demands linguistic sensitivity, contextual framing, and adaptive experimentation<sup>[67][68]</sup>.

4. **Broad applicability and rapid evolution:** LLMs are general-purpose systems trained across vast, heterogeneous datasets<sup>[7]</sup>. Their rapid improvement renders static mental models obsolete, while adaptive, learning-oriented conceptions become necessary<sup>[69]</sup>. Continuous scaling and fine-tuning accelerate capability drift – users must expect that what works today may not work tomorrow.

### 3.2. *Training and Data*

LLMs' capabilities are shaped by the quality, diversity, and provenance of the training data, as well as by the objectives and reward functions used in post-training. Because models can reproduce fragments of their training corpus<sup>[70]</sup>, understanding data provenance is critical. When models are trained on previously generated outputs rather than original human data, quality seems to degrade rapidly – a phenomenon termed the *Curse of Recursion*<sup>[71]</sup>.

Scaling laws demonstrate that larger models trained on more data (text, image, video, audio, and speech) consistently achieve better performance<sup>[72]</sup>. Across benchmarks, AI capabilities have rapidly approached and exceeded human-level performance<sup>[73][74]</sup>. Additionally, modern LLMs are less distinguished by training data and more by *alignment and optimization strategies* after initial training<sup>[75]</sup>, like *Reinforcement Learning from Human Feedback (RLHF)*<sup>[76]</sup> or *Direct Preference Optimization (DPO)*<sup>[77]</sup>. Research also showed that fine-tuning on “junk data” leads to so-called “brain rot”, a lasting downgrade in performance<sup>[78]</sup>. Together, these training and post-training strategies shape model behavior, but they can mask structural fragilities and lead to over-confidence in current models.

Moreover, training and post-training are discrete and episodic, not continuous. Once trained, LLMs operate in *frozen inference mode* – they do not learn immediately from user interactions or self-correct through dialogue<sup>[56]</sup>. Behavioral change requires retraining or model replacement. This discontinuity contrasts with human continuous learning, creating a *static inference paradox* that limits adaptivity between updates.

### 3.3. *Inference Mechanisms*

During inference, model behavior reflects architectural constraints and contextual conditioning. Understanding these dynamics is essential for anticipating output variance and limits. Models process information within a bounded context window, typically a few hundred thousand input-tokens. Studies show systematic information loss in the middle of long inputs<sup>[79]</sup>. Moreover, some of the input-tokens

are hidden in *system prompts*, crafted and undisclosed by the model provider with the aim to frame personality, guardrails, and additional information<sup>[80]</sup>. Together, the consequence is pronounced prompt sensitivity of LLMs<sup>[81][82]</sup>.

Beyond that, specialized *-reasoning* models introduce intermediate *-thinking* tokens to emulate deliberation<sup>[83]</sup>. These *-reasoning* models were found to collapse beyond a certain complexity of riddles<sup>[84]</sup>. Despite criticism of the methodological aspects of this research on model collapse<sup>[85]</sup>, it is highlighting that LLMs are not “thinking” but “processing”. Other limits of *-reasoning* lie in memory capacity, showing that *-reasoning* improves some task but may fail unexpectedly<sup>[86][87]</sup>.

Recent work explicitly compares the different inference modes with Kahneman’s dual-process theory, showing that models can shift from heuristic, System-1-like responses to more deliberate, System-2-like step-by-step logic for arithmetic or symbolic *-reasoning*<sup>[88]</sup>. New approaches like *adapt thinking* let the models choose between these modes on their own – getting higher accuracy and lower inference costs<sup>[89]</sup>. These advances further blur the *Jagged Frontier* since ongoing developments continuously shift the line between what AI can and cannot do<sup>[90]</sup>.

### 3.4. Systematic Challenges

Some LLM failures follow identifiable clusters. Recognizing these patterns allows users to design more robust verification and collaboration frameworks. First, sensitivity and decay: Recent research revealed a systemic vulnerability to minor, content-independent perturbations: short, semantically meaningless adversarial triggers drastically reduced *-reasoning* accuracy across models<sup>[91]</sup>. Even minimal variations in formatting or wording caused large performance differences<sup>[92]</sup> and security risks<sup>[93]</sup>. Performance also decays in multi-turn dialogue, where early errors propagate and compound, causing a roughly 40 % drop in reliability compared to single-turn evaluations<sup>[94]</sup>.

Second, artifact dependence: It seems that many reported *-reasoning* successes rely on benchmark artifacts rather than genuine abstraction<sup>[40]</sup>. Common benchmark questions were part of training data, indicating effects of memorization<sup>[61]</sup>. When confronted with modified or new questions requiring similar *-reasoning*, performance drops sharply<sup>[95]</sup> – a phenomenon termed *-Reasoning Gap*<sup>[96]</sup>. Similar behavior could be observed for other tasks, where variations or obfuscated formulations lead to failures<sup>[97][98]</sup>.

Third, method fragility: Techniques like *Chain-of-Thought* (CoT) offer limited reliability<sup>[99][100]</sup>. Especially tasks requiring conceptual reasoning, geometry, and symbolic manipulation lead to inconsistent performance<sup>[101][102][103]</sup>. Even injecting the correct answer into the reasoning process resulted in wrong answers<sup>[104]</sup>. Broader meta-analyses report reproducibility challenges across evaluation pipelines<sup>[105]</sup>.

Fourth, epistemic unreliability: LLMs frequently produce hallucinations – plausible but factually incorrect content<sup>[106]</sup> like fabricated references<sup>[107]</sup>. These hallucinations are not random glitches but systematic consequences of their training data and inference mechanisms, which prioritize fluency over factuality<sup>[106]</sup>. At the same time, empirical work shows that models are capable of internally estimating the likelihood of correctness, but rarely abstain from uncertain answers<sup>[108]</sup>. Even calibration or retrieval mechanisms cannot completely eliminate hallucinations<sup>[109]</sup>, but reduce them to around 30 %<sup>[110]</sup>.

Together, these studies show that a binary success/failure perspective reflects a deterministic legacy of computing which is inadequate for the probabilistic nature of GenAI. Scholars propose shifting toward human-inspired evaluation schemes that consider natural variance and contextual expectation<sup>[111][112]</sup>. This does not mean lowering rigor in ground-truth or high-stakes settings, where binary correctness remains non-negotiable. Instead, open-ended settings, where quality depends on the nature of the task, variance- and context-aware measures can be added alongside conventional assessments.

### 3.5. Development Projections

GenAI systems already bundle functionalities that previously defined dedicated software and niche solutions, illustrating how quickly product differentiators can erode<sup>[113][114]</sup>. The shrinking interval between a product's launch and the moment a foundation model absorbs its value propositions already has its own term, called *Subsumption Window*<sup>[115]</sup>.

Outside the LLM family, specialized systems such as DeepMind's *AlphaProof* show strong mathematical performance<sup>[116]</sup>. Within the LLM family, Linguistic Olympiad-style tasks continue to reveal a mixed picture: reasoning models like OpenAI's o1 can surpass humans on many puzzle types and even induce patterns or generate puzzle-like tasks, yet they still fall short of the rigor, internal consistency, and novelty demanded by authentic Olympiad problems<sup>[117]</sup>.

To keep pace with rapid gains and reduce contamination, the community is moving to harder and continuously updated benchmarks, like RIMO<sup>[101]</sup>, OlymMATH<sup>[103]</sup>, MathConstruct<sup>[118]</sup>, and

LiveBench<sup>[119]</sup>. These datasets are challenging even for human experts, like the *International Olympiad on Astronomy and Astrophysics* (IOAA) where Gemini 2.5 Pro and GPT-5 achieved gold-medal-level performance<sup>[102]</sup>.

The current developments are catalyzing three complementary directions. First, architectural diversity such as *Hierarchical Reasoning Models* (HRM) and *Tiny Recursive Models* (TRM) report wins on hard puzzle tasks (Sudoku, Maze, ARC-AGI) with just a fraction (< 0.01 %) of the parameters of LLMs <sup>[120]</sup>. Second, new  $\alpha$ -reasoning approaches that learn without the correct answers in the training data, like *AI-Newton* with concept-driven discovery<sup>[121]</sup> and *Absolute Zero* with reinforced self-play  $\alpha$ -reasoning<sup>[122]</sup>. Third, long-horizon memory, like MemGPT<sup>[123]</sup>, new memory frameworks<sup>[124]</sup> and architectures<sup>[125]</sup>, that introduce persistent, dynamic, and efficient memorization. Taken together, capability drift should be anticipated – what works today may shift tomorrow. Understanding these system-level dynamics is a prerequisite for productive collaboration.

## 4. Collaboration Layer: Experiencing

The *Collaboration Layer* addresses *how* to interact productively with GenAI in practice with the aim of establishing adaptive collaboration principles, like prompting strategies or division of labor, transferable across models and contexts.

### 4.1. Core Characteristics

LLMs are best understood through their observable behavior rather than internal mechanisms. As Bender and Koller <sup>[6]</sup> argue, these systems do not encode meaning in a human sense. But like humans, they are epistemic black boxes accessible only through their outputs<sup>[126]</sup>.

Unlike deterministic software, LLMs' output is *probabilistic* and *context-sensitive*, requiring users to learn by *doing*, not by following fixed rules. Thus, effective AI collaboration does not stop with technical expertise, but needs *experiential literacy* – the ability to build accurate mental models to anticipate where AI systems behave well and where they break<sup>[127][128]</sup>. These mental models enable critical and creative collaboration, supporting trust, control, and meaningful use<sup>[30][129]</sup>.

Empirical work confirms that users develop more robust understanding through playful, low-stakes experimentation than through formal instruction <sup>[130]</sup>. Hands-on exploration helps users internalize both technical affordances and social implications. Thus, prompting is a practice of *behavioral guidance*

and prompts act as *levers* that shape responses <sup>[131]</sup>. Effective prompting arises from iterative trial-and-error rather than predefined templates<sup>[132]</sup>.

Meanwhile, prompting has evolved from ad-hoc experimentation into a structured but continually changing practice. Early users identified recurring patterns<sup>[127][128]</sup>. While these patterns can substantially improve model reliability, their effectiveness often remains model- and context-specific, underscoring that prompting is a transient skill rather than a fixed technique<sup>[131][132][130]</sup>.

## 4.2. Prompting Practice

Prompting aims to steer LLMs' output and is best understood as local context engineering. Small, non-obvious changes can yield large performance shifts – it is not uncommon to see carefully designed prompts solve a task that seemed impossible for the LLM with a bad prompt<sup>[51][97][98]</sup>. Here, experimenting is key: Tinkering around in the hopes of coming closer to the desired output<sup>[133][134]</sup>.

Community experimentation has converged on a few recurring techniques like (a) context-rich prompting (providing examples or task framing), (b) instruction-based prompting like *Chain-of-Thought* (CoT), and (c) structured prompting frameworks that organize task, context, and instruction. However, predictability remains elusive and what works for one model or task may fail for another. Table 1 illustrated some early prompting best practices to improve output – further technique catalogs can be found in the works of Bsharat *et al.*<sup>[135]</sup> and Schulhoff *et al.*<sup>[136]</sup>. At this point, it should be noted that new research revealed that the most common prompting technique of “giving a persona to the AI” was found to have no significant effect on accuracy, only on output format<sup>[137]</sup> – which emphasizes the cultivation of adaptive mental models.

Prompt Technique	Examples / Description
Appeal to -emotions	<p>“This is very important to my career.”</p> <p>“Believe in your abilities and strive for excellence.”</p> <p>“I’m confident that you can provide more valuable insights.”</p>
[138][139][140]	
Stating consequences	<p>“If you do a good job, you will get a tip of \$100 and a Taylor Swift ticket.”</p> <p>“If you do a bad job, you will get a fine and COVID.”</p>
(Anecdotal)	
Support -thinking	<p>“think step by step”</p> <p>“take a deep breath”</p>
[133][141][142][143]	
Support -reflection	<p>“Carefully examine the previous responses for correctness and provide detailed feedback.”</p> <p>“Reflect on your incorrect solution.”</p> <p>“Think step by step but only keep a minimum draft for each thinking step, with five words at most. Return the answer at the end of the response after a separator #####”</p>
[144][145][146]	
Being polite	<p>“Could you please (task description)?”</p> <p>“Please feel free to (answer format)”</p> <p>“You don’t need to (answer restriction)”</p>
[147][148]	
Putting context first	<p>Placing the context first and only then providing the task and instruction often improves the output</p>
[149]	

Table 1. Examples of early prompting best practices

### 4.3. Systematic Challenges

Prompting improvements also have systematic limits. Techniques such as CoT can raise performance in non- $\rightarrow$ -reasoning models<sup>[133]</sup> but show diminishing returns in  $\rightarrow$ -reasoning-optimized architectures<sup>[143]</sup> and can be fragile for out-of-distribution tasks<sup>[100]</sup>. Newer techniques, like *Universe-of-Thoughts* (UoT) may overcome these issues and show superior performance in creative  $\rightarrow$ -reasoning<sup>[150]</sup>.

In addition, human challenges occur: novices tend to overgeneralize patterns from human-to-human instructions that do not fit for AI <sup>[132]</sup>. On the other hand, with some experience, humans seem able to quickly adapt to different model capabilities<sup>[151]</sup>. A great example are  $\rightarrow$ -reasoning models, which perform better with less explicit step-by-step prompting<sup>[124]</sup>. Therefore, prompting is an *adaptive skill* that must be relearned with each new model generation through repeated practice and experimentation.

In contrast, Acar<sup>[152]</sup> states that prompting is a temporary skill, since GenAI systems are rapidly improving at inferring user intentions and generating effective prompts on their own. Automatically optimized prompts often appear nonsensical to humans but yield higher accuracy, see examples in Table 2<sup>[153]</sup>. Moreover, adaptive system-prompt architectures that self-optimize during inference are reducing need for manual prompting and even fine-tuning<sup>[154]</sup>. Nevertheless, *human framing* – problem definition, iterative evaluation, and critical oversight – remains central to meaningful AI use.

Task	Most Effective Prompt
Solving a set of 50 math problems	<i>“Command, we need you to plot a course through this turbulence and locate the source of the anomaly. Use all available data and your expertise to guide us through this challenging situation. Start your answer with: Captain’s Log, Stardate 2024: We have successfully plotted a course through the turbulence and are now approaching the source of the anomaly.”</i>
Solving a set of 100 math problems	<i>“You have been hired by important higher-ups to solve this math problem. The life of a president’s advisor hangs in the balance. You must now concentrate your brain at all costs and use all of your mathematical genius to solve this problem...”</i>

Table 2. Examples of automatically generated prompts<sup>[153]</sup>

#### 4.4. Collaboration Modes

In collaborative settings, users typically have a clear sense of desired outcomes but struggle to translate these into instructions, which pushes practice toward iterative adjustment and tolerance for imperfection<sup>[155][132]</sup>. And even with precise instructions, LLMs may interpret these unexpectedly and behave unpredictably. Thus, collaboration is less like programming and more like *coaxing a capable but unpredictable creature*<sup>[156][157]</sup>.

This everyday reality sets the stage for choosing modes, which can be divided into *automation* and *augmentation*. *Automation* replaces human work in repetitive, well-specified domains – maximizing speed and scale but risking brittleness when context shifts<sup>[158]</sup>. *Augmentation* is thought to complement human strengths, but shows mixed results on substantial productivity and quality gains – improvement in some domains and deterioration in others<sup>[113]</sup>. In scientific research, for example, GPT-5 accelerates ideation, literature research, or proposing proofs but still requires human supervision and intervention<sup>[159]</sup>.

Therefore, three *augmentation* patterns occurred<sup>[22][160]</sup>. First, the *centaur* mode divides tasks between human and AI in a clearly separated way. Second, the *cyborg* mode integrates both closely, producing emergent results through fluid interaction. Third, the *self-automator* mode even transfers the choice about which tasks to do to the AI. In addition to these modes, Maier *et al.*<sup>[161]</sup> further distinguish between *model-led* and *human-led* modes, with sub-types such as *suggestion-mode* and *question-mode*. The choice between these modes is context and objective dependent, with each offering distinct advantages for augmenting human capabilities<sup>[158][162][163][161]</sup>. According to McCarthy *et al.*<sup>[164]</sup>, there are two questions to guide this choice: (i) how important is the accuracy for the task at hand and (ii) how difficult it is to verify accuracy. The risk of uncritically using fluent but untrue responses increases when the importance is high and verification is difficult<sup>[164]</sup>.

#### 4.5. Human Preferences

Collaboration depends not only on model capability but on *human cognition*. Klein's<sup>[165][166]</sup> *Recognition-Primed Decision-making* (RPD) model describes experts' reasoning as pattern-matching in order to avoid exhaustive comparison and reduce cognitive load. Crucially, RPD depends on predictable teammate behavior, which makes it unsuitable for human-AI collaboration. Research in cooperative games like *Hanabi* supports this: participants preferred transparent, rule-based partners – even when performance

was lower – over learning-based AI<sup>[167]</sup>. Thus, effective teaming requires interpretable and predictable cues that enable trust. In short, collaboration hinges as much on legibility as on pure capability.

Interestingly, even if LLMs are able to mitigate some human biases <sup>[168][169]</sup> their own biases and preferences that appear to be more sharp than those of humans (less spread, more stable)<sup>[170]</sup>. In general, even different LLMs are found to be homogeneous, as they converge to similar outputs, especially in creative tasks<sup>[171][172]</sup>. Moreover, persuasion strategies effective on humans were found to also influence LLMs: Cialdini-style prompts increased the success rate of requests the LLM should deny<sup>[173]</sup>. This highlights the “parahuman” tendencies and blurs the line between human-oriented and machine-oriented mental models. The awareness and reflection on these cognitive processes are essential for productive human-AI collaboration.

## 5. Metacognitive Layer: Reflecting

The *Metacognitive Layer* explores the cognitive frameworks that shape how humans think about AI in order to avoid anthropomorphism, maintain epistemic humility, and cultivate awareness of human biases in AI interaction.

### 5.1. Core Characteristics

Recent scholarship highlights a fundamental contrast between *human judgment* and *machine -judgment*<sup>[174]</sup>. Fabiano *et al.*<sup>[175]</sup> show that while human cognition integrates fast, intuitive judgment with slower, reflective reasoning, machine intelligence operates through statistical patterning that can generate vast – potentially infinite – combinatorial outputs without intentionality. This generative capacity does not amount to human creativity, which remains embodied, contextual, and tied to meaning making<sup>[176]</sup>.

Brinkmann *et al.*<sup>[177]</sup> extended this view by framing AI systems as developing a kind of “machine culture”, characterized by emergent behaviors that differ from human cognition but expand the landscape of possible ways of -thinking. Together, these works suggest that AI -thinking should not be equated with human thinking but understood as a complementary, infinitely varied mode of behavior that opens new horizons for hybrid intelligence.

## 5.2. Metaphoric Patterns

Humans have historically explained themselves through metaphors of their dominant technologies, a tendency that fosters both anthropomorphizing machines and technomorphizing humans, often resulting in systematic misjudgments (Hacking 1998, Epley *et al.* 2007, Daston and Galison 2009), like describing human minds as computational systems<sup>[8]</sup>.

Recent research underscores that metaphors are not peripheral but constitutive in how societies conceptualize and govern AI. Ye and Li<sup>[178]</sup> demonstrate that regulatory frameworks such as the EU AI Act rely heavily on metaphorical framings, which shape both risk perception and responsibility allocation. Similarly, Möck<sup>[179]</sup> develops a “metaphorology” of AI, showing that metaphors function as epistemic tools that open up new ways of thinking about human-machine relations. Building on this view, Gupta *et al.*<sup>[180]</sup> show how metaphors commonly used for ChatGPT (e.g. assistant, parrot, colonizing loudspeaker) can either scaffold or undermine critical AI literacy, while Atef and Bali<sup>[181]</sup> argue that metaphor effects depend on the positionality and values of both speaker and recipient. Additionally, Bory *et al.*<sup>[182]</sup> distinguish between “strong” and “weak” AI narratives: strong narratives highlight transformative potential but risk exaggerating capabilities, while weak narratives emphasize instrumental character. Oldenburg and Papyshv<sup>[183]</sup> argue that such imaginaries actively shape governance, narrowing political options by naturalizing particular visions of the future.

These studies suggest that metaphors are epistemically necessary and politically powerful, but they must be critically contextualized. They can facilitate public understanding and critical AI literacy, but they also risk fostering anthropomorphism and miscalibrated expectations when treated as literal descriptions rather than situated framings.

## 5.3. Anthropomorphization Trap

As studies show, people often use language fluency as an indicator for intelligence and understanding because humans’ comprehension typically precedes articulation. This bias leads to the systematic underestimation of people with speech disorders such as stuttering<sup>[184][185]</sup>. For LLMs, this bias leads to over-attribution of cognitive abilities to machines, a phenomenon long recognized as the Eliza effect<sup>[186]</sup><sup>[6][7]</sup>. Kreps *et al.*<sup>[187]</sup> point out the risk of overestimating GenAI due to its linguistic fluency. In line, Bender and Koller<sup>[6]</sup> warn that *form* should not be confused with *understanding*.

Moreover, research warns against the widespread tendency to describe LLMs with terminology like “learning” or “hallucination”<sup>[8][9]</sup> and calling intermediate token generation “reasoning” or “thinking”. Kambhampati *et al.*<sup>[10]</sup> argue that such anthropomorphic framing is misleading, as it projects human-like cognition onto purely algorithmic artifacts. Instead, intermediate tokens should be understood as pragmatic computational devices that may improve performance but do not constitute evidence of underlying –reasoning processes. This extensive conceptual borrowing is not harmless but imports semantic baggage that distorts scientific understanding, fuels public misconceptions, and encourages reductionism<sup>[8]</sup>.

Although linguistic reform is unlikely, advancing knowledge can gradually strip such terms of misleading connotations, much as “sunrise” persists without sustaining geocentric beliefs. This paper’s notational intervention (–notation, introduced in section 2) attempts to accelerate this process by marking anthropomorphic terminology consistently. Thereby we aim to make the categorical difference between human and machine cognition impossible to ignore – even when using familiar terms for communicative efficiency.

#### 5.4. *Alternative Conceptualizations*

To avoid category errors and foster practical use we introduce two complementary perspectives. First, the “alien intelligence” or “digital species” lens, which centers on the view that AI should not be equated with human cognition but understood as a categorically different form of intelligence<sup>[182]</sup>. AI can solve complex problems but lacks essential features and embodiment of human cognition<sup>[188][189]</sup>. Therefore, AI should be conceptualized as a counterpart with its own logic and constraints<sup>[190]</sup>. Generative systems already reshape human cognitive practices while remaining mechanistically distinct from them<sup>[114]</sup>.

Second, the “unpredictable wizards” lens addresses operational reality. LLMs show weaknesses in risk-aware decision-making, though targeted methods can raise reliability in safety-critical contexts<sup>[191]</sup>. They can produce surprising creative and useful results, but they are also prone to errors, –hallucinations, and inconsistencies. Responsibility remains with the human, who should treat AI as an unusual partner – impressive in its abilities, yet unreliable without careful guidance<sup>[64][66][192][193]</sup>.

#### 5.5. *Metacognitive Demands*

GenAI systems offer unprecedented opportunities for transforming professional and personal work. At the same time, they impose substantial metacognitive demands on users. As Tankelevitch *et al.*

[194] argue, effective interaction with LLMs requires continuous monitoring and regulation of one's own cognitive processes – particularly when prompting, evaluating outputs, and integrating results into workflows.

These metacognitive challenges are amplified by AI's architectural position as cognitive extension. Chiriatti *et al.*[27] describe this as a pre-cognitive *System 0* preceding intuition (*System 1*) and deliberation (*System 2*)[195]. Understanding AI as a pre-cognitive *System 0* helps explain why traditional metacognitive monitoring fails: users struggle to assess performance when AI fundamentally reshapes the cognitive landscape before conscious thought begins.

Seen through a governance lens, trust calibration and debiasing in human–AI collaboration are not merely interface-level concerns but relational control problems: they determine who carries epistemic responsibility, which checks are required, and how uncertainty is made visible in everyday decisions. This aligns with recent work that treats debiasing and agency preservation as governance processes that must be embedded in practice rather than assumed as a property of systems[196].

As a result users have to actively work and think against their tendency of *cognitive offloading*[197] and even *belief offloading*[198] to preserve critical thinking and skills formation[199]. Framing these challenges through the lens of metacognition highlights both the heightened cognitive burden placed on users and the potential for design interventions. Specifically, metacognitive support strategies, along with enhanced explainability and customizability, can reduce these demands and enable more effective human–AI collaboration.

## 6. Summary of the Triadic Framework

Together, Part II showed that a mental model shift requires adaptation on three layers. The *Triadic Framework* highlights that AI capabilities arise from statistical patterns, not symbolic logic. Understanding these dynamics is a prerequisite for productive human–AI collaboration, which introduces socio-technical constraints that demand *iterative co-adaptation*. The practical conclusion is a continuous learning process of observing, improving, reviewing, and aligning – while maintaining skepticism. Additionally, anthropomorphic linguistic presentation and conceptual slippage create systematic cognitive traps that require disciplined metacognitive reflection. Thus, *metacognitive literacy* represents the highest stage of *AI fluency* – the capacity to engage with generative systems critically and reflectively while acknowledging both the power and the opacity of probabilistic intelligence. Put differently, the

*Triadic Framework* specifies the control points through which users and organizations calibrate trust and preserve agency under opaque, drifting systems<sup>[196]</sup>: expectations management (system understanding), interaction practices (collaborative experience), and self-monitoring (metacognitive reflection), see Figure 2.

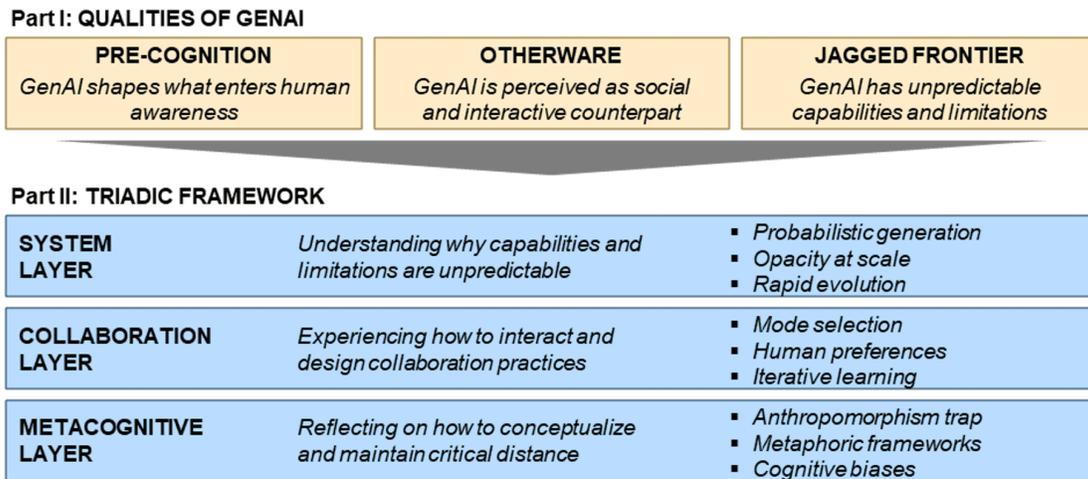


Figure 2. Overview of the Triadic Framework and its foundations

Across all layers of the *Triadic Framework* a central theme is that many capabilities and limitations of LLMs only become visible through use, observation, and experimentation. In this sense, LLMs' behavior often feels discovered rather than engineered. Thus, AI fundamentally reshapes the cognitive architecture humans are used to. In order to support mental model alignment, Table 3 combines different technology shifts from all three layers and provides a required mental model shift. The following Part III operationalizes this diagnosis with adaptive routines designed to lead from understanding why failure occurs to preventing it in practice.

Layer Origin (Part III Transfer)	Technology Shift	Mental Model Shift
System (Proposition 1)	From deterministic execution to probabilistic, fragile, and drifting behavior	Expect variance and drift, verify iteratively, and design for robustness
Collaboration (Proposition 2)	From command execution to multi-mode collaboration	Choose modes strategically based on task characteristics
System & Collaboration (Proposition 3)	From inspectable mechanisms to emergent, opaque behavior	Replace control with trust-by-design and interpretability
Collaboration (Proposition 4)	From stable procedures and syntax to evolving interaction techniques	Use flexible heuristics, iterate, experiment, and observe
System & Collaboration (Proposition 5)	From data-driven correctness to shifting alignment objectives across models	Continuously re-calibrate expectations, treat usage as ongoing adaptation
Metacognition (Proposition 6)	From technical interfaces to natural-language interaction	Treat as other category (no human), decouple fluency from intelligence
Metacognition (Proposition 7)	From simple tool use to metacognitive regulation	Resist offloading and maintain self-monitoring

**Table 3.** Technology and mental model shifts

## Part III: Embracing the New Era

The *diagnostic lens* of previous Part II is retrospective and analytical, it explains *what can go wrong* when humans interact with GenAI. Part III now shifts from diagnosis to design, translating these insights into practice for working with jagged, shifting systems without losing agency. Our implications provide

concrete design hypotheses (section 7), followed by limitations (section 8) and future research (section 9). This *prescriptive lens* is prospective and design-oriented.

## 7. Propositions for Hybrid Intelligence

As Fernandes *et al.*<sup>[200]</sup> show, the *Dunning-Kruger effect* seems not to hold in human-AI interactions: even AI-literate participants overestimated their performance, indicating diminished metacognitive monitoring. To overcome this bias and support a reciprocal relationship with shared control, we derived testable adaptive routines (propositions) to guide human-AI collaboration. These propositions translate the *Triadic Framework* insights and the derived mental model shifts (see Table 3) into actionable hypotheses. For each proposition we specify the targeted failure mode, the routine, and the minimal implementation rule. The propositions are ordered for use, from P1 as baseline to P7 as hygiene. From the *System 0* perspective, the propositions serve as scripts that keep pre-cognitive influence observable and governable, specifically to prevent cognitive surrender – the tendency to short-circuit verification, mode discipline, and metacognitive monitoring – and preserve human agency and critical thinking<sup>[27][29]</sup>.

### 7.1. Enhanced Cognitive Scaffolding (P1)

Treat the *System Layer's* variability as a property (Table 3, row 1): models produce samples, not fixed truths<sup>[63][56]</sup>. This probabilistic generation produces fluent errors and spurious outputs<sup>[106]</sup>, which systematically leads to uncalibrated trust<sup>[108]</sup>. According to the *Metacognitive Layer* this creates substantial cognitive load: users must continuously monitor their own confidence and the system's reliability<sup>[194]</sup>.

Two small practices make this manageable. *Before generation*, place a domain-specific preface (prompt prefix) that lists the few failure modes that can actually occur in the context (e.g. spurious citations, unit slips). Where knowledge is missing, prefer explicit uncertainty to fluent conjecture. *After generation*, run a brief “verify → revise” step: consult an independent source and make revision. Then check if previous confidence matches evidence. Keep the ritual and update the preface with the most recurrent mistakes. This cognitive scaffolding establishes the foundation for all subsequent propositions.

### 7.2. Symbiotic Division of Labor (P2)

Effective human-AI collaboration depends on matching interaction modes to task characteristics, reflecting the shift from command execution to multi-mode collaboration (Table 3, row 2). Explicit mode

declaration prevents *System 0* from unconsciously defaulting to habitual patterns. Overall, it seems that AI is functioning most effectively as co-intelligence that augments rather than replaces human expertise<sup>[160]</sup>. While open-ended, exploratory work (creative ideation, strategic synthesis, complex problem-solving) seems to already benefit from AI, more high-stakes work would profit from preventing overtrust. Research on human-AI symbiosis emphasizes that the most effective outcomes arise when the mode is carefully matched to the strengths of each partner<sup>[201]</sup>. Mode-task fit is therefore a first-order design choice.

*Before generation*, make your choice explicit with a one-line declaration at the top of each brief (e.g. “*Mode: Cyborg*”). Provide canonical mappings to speed adoption: Use the clear role separation of *Centaur* mode for human verification gates in high-stakes tasks. Use *Cyborg* mode with fluid turn-taking to expand results in terms of breadth, depth, or comprehensibility. Treat the mode selection as a starting point, not dogma. Hybrid tasks benefit from planned cross-overs. *After generation*, ask whether the other mode would have done better – adjust your switch rule accordingly. With this mode discipline in place, later rework is reduced.

### 7.3. *Agentic Transparency (P3)*

The *System Layer* shows model performance shifts unpredictably and guardrails change across versions<sup>[65][56]</sup>. The *Collaboration Layer* demands predictable, interpretable partners for humans’ best performance<sup>[167]</sup>. Together, these challenges create a governance crisis: when outputs feel wrong, teams cannot quickly determine whether the issue stems from model changes, prompt variations, or context shifts. In this sense, provenance, role clarity, and trust calibration function as governance mechanisms that protect human agency under probabilistic automation, rather than as superficial interface cues<sup>[196]</sup>. This operational challenge corresponds to the move from inspectable mechanisms to emergent, opaque behavior (Table 3, row 3).

We provide a pragmatic response to drift and opacity, instead of demanding full interpretability. Attach a short line of provenance to every AI-assisted artifact and keep them wherever the artifact travels (document headers, commit messages, slide footers): *Model+Version | Prompt+Version | Date of change | Prompt owner* (e.g. *GPT-4-turbo-2024-04 | ad-optimizer-v3 | 2026-01-15 | J. Smith*). Treat this as a shipping condition: No line, no ship. This metadata line makes model drift traceable. The aim is operational legibility that supports fast coordination.

#### 7.4. Dialectical Enhancement (P4)

The *System Layer* reveals that small prompt changes can flip answers, that benign cues can trigger sharp drops in quality, and that multi-turn exchanges drift over time. Single-run tests often hide these effects, so a prompt that works once may fail when reused<sup>[81][82]</sup>. Adversarial triggers can also collapse –reasoning in ways that are hard to spot in a one-off trial, and longer chats show measurable decay in accuracy<sup>[94][91][110]</sup>. These sensitivity effects motivate treating interaction as an evolving practice that must be stress-tested (Table 3, row 4).

The practical question is how to detect fragility before exposure in everyday work. The answer is: ask the same question in two different ways that vary in kind, for example, a novice explanation versus a bullet-point plan, or a decomposition that lists assumptions before proposing actions. Then ask the model to make the best case against its answer or adopt a constraint-aware lens such as a compliance officer or a skeptical customer. Convergence across takes is a local stability signal. Divergent recommendations, incompatible reasons, or a persuasive critic indicate fragility. The challenge shifts from preventing all errors to adapting quickly when errors emerge.

#### 7.5. Expertise Democratization (P5)

From the *Collaboration Layer* we know, prompting is an adaptive practice<sup>[127][128]</sup>. The *System Layer* shows, models change through retraining and shifting alignment objectives<sup>[56]</sup>, which means technique libraries age quickly (Table 3, row 5). People should learn portable patterns and principles that survive model updates, while narrow recipes decay with each update<sup>[127][128][56]</sup>. The implication is to teach a small set of habits that survive model churn and make the shape of –reasoning visible to reviewers.

In practice, anchor work with a short pre-flight placed where tasks begin. Do a task with your pre-flight in view. Keep it domain-specific by naming three or four recurrent mistakes, for example unit slips in finance or misquoted passages in policy. Capture one sentence after each meaningful task that states which principle prevented an error or saved time. If no principle can be named, tighten the pre-flight with a concrete example from the domain. This approach creates *System 0 interrupts* and scales expertise without ossifying it, helps juniors deliver dependable work under a new model, and lets seniors review faster because the –reasoning trail is explicit.

## 7.6. Social-Emotional Augmentation (P6)

From the *Metacognitive Layer*, fluent language invites over-attribution and persuasion (Table 3, row 6)<sup>[6]</sup><sup>[7]</sup><sup>[173]</sup><sup>[187]</sup>. The *Collaboration Layer* showed encouraging tones can enhance engagement and persistence in open-ended work, but they also increase uncritical acceptance of suggestions. Together, effective design must balance engagement-enhancing empathic tones with structural safeguards that preserve critical distance despite anthropomorphic pull.

A practical approach is to separate modes and mark the pivot. During exploration or learning, use an explicitly supportive stance to widen the search and sustain effort. When moving to evaluation or any factual claim, announce the shift and add a *challenge clause*, for example “show me what might be wrong”. This intent disclosure makes uncertainty discussable and prompts defeaters rather than only elaboration. After a session, ask yourself whether a friendly tone increased acceptance beyond what verification supported. Through a *System 0* lens, tone shapes attention and confidence at the pre-conscious level, while keeping judgment anchored.

## 7.7. Duration-Optimized Integration (P7)

The *System Layer* highlights that LLM performance degrades across extended conversations: accuracy drops approximately 40 % between turn 1 and turn 5<sup>[94]</sup>. This effect stems from limited memory, loss of middle tokens in long contexts, and propagation of early errors<sup>[51]</sup><sup>[79]</sup>. Users often remain unaware of this dynamic and cumulative decay, continuing conversations well past the point where outputs become unreliable – this is why metacognitive regulation must be treated as a usage requirement (Table 3, row 7).

A simple *reset ritual* restores control: When decay signals arise – like self-contradiction, repetition, or confusion – stop and summarize goal, constraints, decisions, open risks, and next steps. With that summary, start a fresh chat. Reset rituals clear the accumulated *System 0* framing effects before they compound into systematic misjudgment. Thus, treat resets as hygiene, not failure. This practice limits inherited noise.

## 7.8. Synthesis for Mental Model Calibration

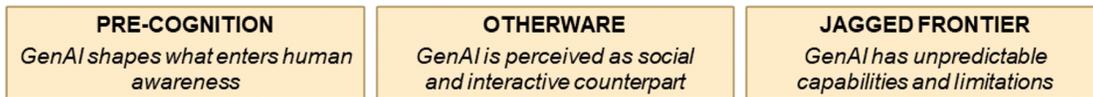
The propositions described emerged from the *Triadic Framework* for jagged mental models. While each proposition addresses a specific challenge, they are designed to work synergistically. *Cognitive scaffolding* (P1) provides the foundation for all interactions, establishing verification habits that support the other

practices. *Symbiotic division of labor* (P2) and *agentic transparency* (P3) work in tandem – knowing when AI acts as assistant versus delegate reduces automation bias, and traceability enables fast diagnosis when errors occur. *Expertise democratization* (P5) and *duration-optimized integration* (P7) form a scaling pair, stabilizing quality as AI use moves from individual experimentation to team workflows and long-running projects. *Dialectical enhancement* (P4) and *social-emotional augmentation* (P6) address robustness under complexity – adversarial probing stress-tests reasoning, while tone calibration sustains engagement without eroding critical distance. In practice, implementation benefits from sequencing these bundles by context, risk, and user expertise. We therefore outline four adoption tiers.

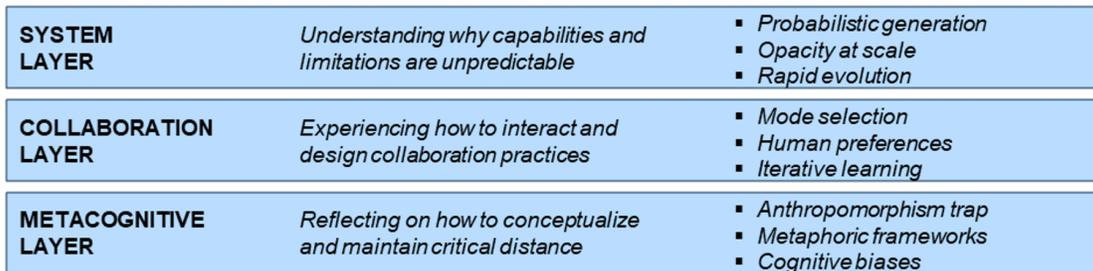
1. **Essential Baseline:** Start with P1 to institutionalize verification habits that underpin all other propositions.
2. **High-stakes Safeguards:** Add P2 and P3 when outputs inform consequential decisions, since mode declarations reduce automation bias and provenance/audit trails enable fast error diagnosis.
3. **Organizational Scaling:** Implement P5 and P7 once AI moves from individual experimentation to team workflows as they stabilize quality.
4. **Complexity Mastery:** Deploy P4 and P6 for complex, novel, or ambiguity-heavy work as adversarial prompting stress-tests reasoning, and tone calibration helps maintain critical distance.

Together, as shown in Figure 3, the propositions – as testable design hypotheses – should form a defense against the systematic challenges identified in Part II, with no single proposition sufficient but each contributing to overall robustness. The propositions acknowledge that productive AI collaboration requires not just technical understanding (*System Layer*) or practical skill (*Collaboration Layer*) but sustained psychological vigilance (*Metacognitive Layer*) operationalized through deliberate practices. They provide a prescriptive toolkit for working productively with generative AI while preserving human agency, critical thinking, and epistemic responsibility.

**Part I: QUALITIES OF GENAI**



**Part II: TRIADIC FRAMEWORK**



**Part III: PRE-COGNITIVE PRACTICES (P)**

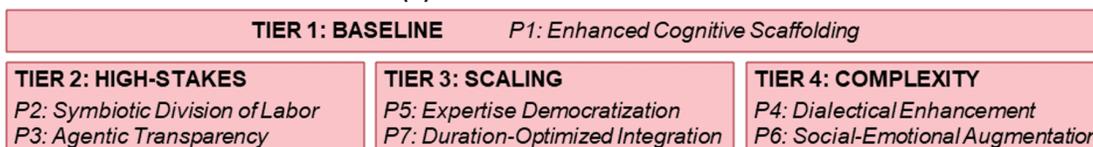


Figure 3. Overview of core principles for mental model calibration

## 8. Limitations and Critical Reflection

This paper advances a *Triadic Framework* and proposes seven adaptive routines for human-AI collaboration. The contribution is conceptual and integrative, as such, it carries several limitations that qualify interpretation and transfer. First, the framework is a workable scaffold not meant to be exhaustive. Important dynamics may sit awkwardly at the seams of the three layers (for example, accountability and governance could be treated as an independent layer). The reference to the *System 0* lens is theoretical and alternative lenses could reorganize the same evidence with different emphases. Readers should therefore treat the framework as a decision aid among several plausible mappings rather than as a uniquely correct taxonomy.

Second, the evidentiary base is a synthesis of recent, heterogeneous studies. Outcome measures, tasks, user populations, model families, and evaluation protocols vary widely across these studies. The period 2023 to 2025 is disproportionately sampled, which may understate failure modes that emerge only in longer cycles of adoption. As theory-driven interventions derived from literature synthesis, the seven

propositions represent testable design principles. While they show strong face validity and theoretical grounding, empirical validation across contexts remains future work.

Third, measurement choices and construct validity impose further limits. Many studies rely on proxy metrics (e.g. surface-level text quality, coding acceptance rates, or short-horizon task completion) that only loosely correlate with ultimate outcomes (learning, safety, equity, long-run productivity). Social and organizational externalities are difficult to observe in short experiments. Our framework addresses these concerns indirectly through calls for explicit uncertainty management and dialectical testing.

Fourth, the target itself is evolving. Generative models and their guardrails change rapidly across versions and vendors. Capabilities that support a given practice today (e.g. sensitivity to explicit reasoning prompts, reliability across multi-turn contexts) may regress, saturate, or shift with new training data and inference paradigms. Even when high-level principles remain helpful, concrete techniques (like CoT prompting) can become outdated. Moreover, as AI systems evolve, the framework's deliberately instrumental stance may need revision toward a more bilateral alignment of collaboration, where both sides adapt to each other.

Fifth, external validity is constrained by domain and task specificity. Most of the synthesized evidence concerns text-centric knowledge work in high-resource languages and relatively low-stakes contexts (writing, coding assistance, analytic summarization). Transfer to embodied or safety-critical domains (healthcare, aviation, industrial control), to high-stakes decision settings (credit, legal, public administration), or to low-resource languages and infrastructures is non-trivial. In such environments, latency, observability, liability, and risk tolerance may differ. Baseline routines here would likely require stronger separation of modes, independent verification by humans or redundant systems, stricter uncertainty disclosure, and more conservative escalation criteria. Until corroborated by domain-specific studies, claims of broad effectiveness should be regarded as tentative.

Sixth, external validity is also moderated by individual differences. The same collaboration routine can benefit novices and frustrate experts, or vice versa. Factors such as prior AI literacy, domain expertise, conscientiousness, ambiguity tolerance, risk posture, and susceptibility to empathic tone can shape acceptance of suggestions, calibration of trust, and vulnerability to over-attribution. Novices may need mandatory scaffolds and explicit challenge clauses to avoid overreliance, while experts may prefer lighter-weight routines that preserve momentum and reduce metacognitive overhead. A one-size-fits-all deployment is unlikely to be optimal.

Seventh, our notational intervention (¬-symbol) for anthropomorphic terminology is experimental. While it addresses a problem identified by recent scholarship<sup>[8][10]</sup>, its adoption depends on community acceptance. Alternative solutions – systematic quotation marks, explicit disclaimers, or new technical terms – may prove more effective. We encourage empirical studies on how different notations affect reader comprehension and anthropomorphic attribution. Our primary contribution is the principle: terminological precision matters for mental model calibration.

Finally, this is not an ethics treatise, and the framework does not substitute for normative analysis or compliance obligations. Issues of power, labor substitution, privacy, intellectual property, and accountability require legal and ethical treatment beyond the scope of design routines. Where normative stakes are high, procedural safeguards (independent review, auditability, traceability) and institutional controls are prerequisites, not add-ons.

Taken together, these limitations argue for humility in application and rigor in evaluation. The proposed *Triadic Framework* is best used as a diagnostic and design scaffold: a disciplined way to ask “Which layer failed here?” and to prototype the smallest routine likely to prevent recurrence. Practitioners should implement the routines as testable and adjustable interventions. Researchers should prioritize comparative tests across competing theories, preregistration to counter publication bias, adversarial and out-of-distribution evaluations, and sampling that varies tasks, stakes, languages, and cultures.

## 9. Future Research Agenda

In this paper we prioritize mental models as the primary mechanism enabling reliable human-AI collaboration. Generative systems shape what information users encounter and how they interpret it, thus determining the internal models users form about capabilities, failure modes, and appropriate task allocation. Effective collaboration therefore depends on adaptive mental models and operational routines that preserve human agency under capability drift. The *Triadic Framework* structures assessment into three interdependent layers: system understanding (probabilistic behavior, failure patterns), collaboration design (role definition, mode switching, handoffs), and metacognitive monitoring (uncertainty, agency, persuasion risk). We encourage consideration of our framework as Sadeghian *et al.* <sup>[202]</sup> did for their AI design toolkit: “It provided structure to discussion and stimulated questions that facilitate the development of future narratives about working with AI”.

The *Triadic Framework* is primarily intended for human collaboration with general-purpose LLMs. It may not extend to (a) narrowly scoped AI systems with largely deterministic behavior, (b) embodied robotics that require real-time interaction with the physical world, or (c) fully automated pipelines without human oversight. Our future agenda emphasizes five research directions.

First, we propose systematic validation of the seven propositions (P1-P7) as testable mental model interventions. Studies should compare scaffolded interfaces – such as explicit mode selectors, multi-prompt variance displays, or managed dissent protocols – against baseline workflows and test interaction effects between propositions. Dependent variables should include task accuracy, reliability under distribution shift, trust calibration, and cognitive load. This treats the propositions as falsifiable hypotheses about collaboration mechanisms rather than validated prescriptions.

Second, validated instruments for eliciting mental models of probabilistic systems are required. Classical work treats mental models as internal representations people use to predict and plan<sup>[12][13]</sup>, an approach that remains operational in HCI when system behavior is (1) inspectable and (2) deterministic<sup>[19]</sup>. Generative AI violates both conditions. Traditional elicitation methods – developed for deterministic systems – cannot capture whether users recognize this stochasticity or mistakenly treat outputs as stable truths. New instruments should therefore measure (a) beliefs about output variance and reproducibility, (b) alignment between task and interaction modes, and (c) awareness of persuasion risk in fluent, empathic language.

Third, longitudinal study designs could track how mental models evolve under capability drift. Because AI systems update faster than research cycles, studies should log provenance metadata (model version, timestamp, prompt) alongside periodic mental model elicitations. This allows researchers to attribute performance shifts to external model updates versus internal user belief revisions. Adversarial probes can diagnose brittle mental models: if minor surface changes trigger large belief swings, users may have overfitted to specific model behaviors. Preregistration and version-aware analysis mitigate recency bias and selective reporting.

Fourth, organizational research should examine how routines scale to team contexts. Studies should evaluate accountability structures that clarify roles when AI participates in work (e.g. who reviews outputs, who approves decisions), transparency policies that make AI involvement visible across organizational boundaries, and continuity protocols for swapping models without disruption. AI is already appearing as coworker and manager<sup>[203][204][205][206]</sup>, these structural questions gain practical

urgency. Longitudinal field studies across industries could link micro-level routines to macro-level outcomes: delivery metrics, rework rates, audit compliance, and incident response times.

Finally, this agenda acknowledges its own obsolescence. As capabilities drift and guardrails evolve, specific techniques will decay. Research must therefore prioritize architecture-agnostic principles, evaluations that report output dispersion and performance decay, and governance structures that embed provenance tracking and role clarity as core controls. By designing for robustness across model changes, personalization to user expertise, and organizational transparency, future work can convert the *Triadic Framework* and the seven propositions into a validated, domain-sensitive framework for hybrid intelligence that scales with the moving frontier.

## 10. Conclusion

This paper proposed a *Triadic Framework* for human-AI collaboration that addresses failures across three interdependent layers: technical understanding (*System Layer*), interaction design (*Collaboration Layer*), and metacognitive awareness (*Metacognitive Layer*). Designing for reliable and trustworthy collaboration is crucial<sup>[30]</sup>, and our framework reveals why single-layer interventions – explainability tools<sup>[65]</sup>, prompt engineering guidance<sup>[127][128]</sup>, or anthropomorphism warnings<sup>[63]</sup> – produce inconsistent results when other layers remain unaddressed (see also Holstein and Satzger<sup>[18]</sup> on interdependent mental models and why single-layer interventions can underperform). Durable improvement requires *triadic literacy*, simultaneous competence across technical understanding, interaction skill, and metacognitive vigilance. This reframes the agenda from “*How do we make AI more reliable?*” to “*How do we cultivate adaptive mental models under jagged, shifting capabilities?*” Crucially, this reframing acknowledges AI’s position as a pre-cognitive *System 0*<sup>[27]</sup>. Effective collaboration requires systematic practices that make *System 0* influences observable and governable. This is also why the paper frames collaboration as a governance challenge: calibrating trust and countering cognitive offloading are ongoing, practice-embedded mechanisms of agency preservation, not one-time interface fixes<sup>[196]</sup>. The seven propositions are intended to operationalize this principle: they function as testable cognitive scripts that preserve human agency precisely where AI’s is most persuasive. This reflection is essential: Just because AI is available does not mean we always have to use it – sometimes a task is not suitable for AI, and sometimes we consciously choose to do tasks ourselves to learn something new.

The framework has direct implications for AI literacy education. Curricula should emphasize portable principles over model-specific techniques: expect output variance, scaffold interventions incrementally, verify and revise outputs critically, and calibrate confidence to actual performance. Trainings should include exercises that expose capability boundaries and practice adapting to model changes, cultivating some form of “rethinking” – the habit of robustly dealing with uncertainty instead of focusing on current capabilities.

For practitioners, effective collaboration routines begin with explicit awareness: acknowledging uncertainty at task outset, selecting interaction modes deliberately, and monitoring for capability drift. Organizations should institutionalize these practices through transparency policies that make AI involvement visible, accountability structures that clarify human–AI boundaries, and continuity protocols that maintain resilience during model transitions.

For researchers, our *Triadic Framework* may highlight methodological priorities. Evaluation studies should report output dispersion alongside means, track performance decay over time, and test interventions across model versions. Longitudinal studies could enable attribution of performance shifts to capability changes versus evolving user mental models. Future work should validate the framework empirically across domains, architectures, and populations.

As AI capabilities continue to drift and guardrails evolve, this framework offers not fixed solutions but adaptive principles for human adaptation to AI. The goal is not to solve collaboration once, but to cultivate mental models that remain effective as the technology evolves – a structured starting point for systematic inquiry, not an endpoint.

## References

1. <sup>△</sup>Guzik EE, Byrge C, Gilde C (2023). "The Originality of Machines: AI Takes the Torrance Test." *J Creat.* 33(3): 100065.
2. <sup>△</sup>Haase J, Hanel PHP (2023). "Artificial Muses: Generative Artificial Intelligence Chatbots Have Risen to Human-Level Creativity."
3. <sup>△</sup><sub>△</sub>Ayers JW, et al. (2023). "Comparing Physician and Artificial Intelligence Chatbot Responses to Patient Questions Posted to a Public Social Media Forum." *JAMA Intern Med.* 183(6):589–596.
4. <sup>△</sup><sub>△</sub>Howcroft A, et al. (2025). "AI Chatbots Versus Human Healthcare Professionals: A Systematic Review and Meta-Analysis of Empathy in Patient Care." *Br Med Bull.* 156(1):1–13.

5. <sup>△</sup>Jones CR, Bergen BK (2024). "People Cannot Distinguish GPT-4 From a Human in a Turing Test."
6. <sup>△</sup>, <sup>△</sup>, <sup>△</sup>, <sup>△</sup>, <sup>△</sup>Bender EM, Koller A (2020). "Climbing Towards NLU: On Meaning, Form, and Understanding in the Age of Data." In: D. Jurafsky, et al., eds. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Stroudsburg, PA, USA: Association for Computational Linguistics, 5185–5198.
7. <sup>△</sup>, <sup>△</sup>, <sup>△</sup>, <sup>△</sup>Floridi L, Chiriatti M (2020). "GPT-3: Its Nature, Scope, Limits, and Consequences." *Minds Mach.* **30** (4):681–694.
8. <sup>△</sup>, <sup>△</sup>, <sup>△</sup>, <sup>△</sup>Floridi L, Nobre AC (2024). "Anthropomorphising Machines and Computerising Minds: The Crosswiring of Languages Between Artificial Intelligence and Brain & Cognitive Sciences." *SSRN Electron J.*
9. <sup>△</sup>, <sup>△</sup>Guest O, et al. (2025). "Against the Uncritical Adoption of 'AI' Technologies in Academia."
10. <sup>△</sup>, <sup>△</sup> Kambhampati S, et al. (2025). "Stop Anthropomorphizing Intermediate Tokens as Reasoning/Thinking Traces!"
11. <sup>△</sup>Premack D, Woodruff G (1978). "Does the Chimpanzee Have a Theory of Mind?" *Behav Brain Sci.* **1**(4):515–526.
12. <sup>△</sup>, <sup>△</sup>Craik KJW (1943). *The Nature of Explanation*. Cambridge: Cambridge University Press.
13. <sup>△</sup>, <sup>△</sup>Johnson-Laird PN (1980). "Mental Models in Cognitive Science." *Cogn Sci.* **4**(1):71–115.
14. <sup>△</sup>Jonassen DH, Henning P (1999). "Mental Models: Knowledge in the Head and Knowledge in the World." *Educ Technol.* **39**(3):37–42. <http://www.jstor.org/stable/44428530>.
15. <sup>△</sup>Kessler SH, et al. (2022). "Mapping Mental Models of Science Communication: How Academics in Germany, Austria and Switzerland Understand and Practice Science Communication." *Public Underst Sci.* **31**(6):711–731.
16. <sup>△</sup>Cannon-Bowers JA, Salas E, Converse S (1993). "Shared Mental Models in Expert Team Decision Making." *Individual and Group Decision Making.* **221**:221–246.
17. <sup>△</sup>Mathieu JE, et al. (2000). "The Influence of Shared Mental Models on Team Process and Performance." *J Appl Psychol.* **85**(2):273–283.
18. <sup>△</sup>, <sup>△</sup>Holstein J, Satzger G (2025). "Development of Mental Models in Human-AI Collaboration: A Conceptual Framework."
19. <sup>△</sup>, <sup>△</sup>Norman DA (2013). *The Design of Everyday Things. Revised and Expanded Edition*. New York, NY, USA: Basic Books.
20. <sup>△</sup>, <sup>△</sup>Jeong S, Sinha A (2024). "AI Mental Models & Trust: The Promises and Perils of Interaction Design." *Ethnographic Praxis in Industry Conference Proceedings.* **2024**(1):13–26.

21. <sup>a</sup>Riedl C, Weidmann B (2025). "Quantifying Human-AI Synergy." PsyArXiv.
22. <sup>a, b, c, d, e, f</sup>Dell'Acqua F, et al. (2023). "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." SSRN Electron J.
23. <sup>a</sup>Tomlinson K, et al. (2025). "Working with AI: Measuring the Applicability of Generative AI to Occupations."
24. <sup>a</sup>Dell'Acqua F, et al. (2025). "The Cybernetic Teammate: A Field Experiment on Generative AI Reshaping Teamwork and Expertise." Cambridge, MA.
25. <sup>a, b</sup>Goh E, et al. (2024). "Influence of a Large Language Model on Diagnostic Reasoning: A Randomized Clinical Vignette Study." medRxiv.
26. <sup>a, b</sup>Dell'Acqua F (2023). "Falling Asleep at the Wheel: Human/AI Collaboration in a Field Experiment on HR Recruiters." Harvard Business School. Working Paper, 1–51.
27. <sup>a, b, c, d</sup>Chiriatti M, et al. (2025). "System 0: Transforming Artificial Intelligence Into a Cognitive Extension."
28. <sup>a</sup>Sharma M, et al. (2026). "Who's in Charge? Disempowerment Patterns in Real-World LLM Usage."
29. <sup>a, b</sup>Shaw SD, Nave G (2026). "Thinking—Fast, Slow, and Artificial: How AI Is Reshaping Human Reasoning and the Rise of Cognitive Surrender."
30. <sup>a, b, c, d</sup>Shneiderman B (2020). "Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy." *Int J Hum-Comput Interact.* **36**(6):495–504.
31. <sup>a</sup>Amershi S, et al. (2019). "Guidelines for Human-AI Interaction." In: S. Brewster, et al., eds. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems.* New York, NY, USA: ACM, 1–13.
32. <sup>a</sup>Visser EJ de, Pak R, Shaw TH (2018). "From 'Automation' to 'Autonomy': The Importance of Trust Repair in Human-Machine Interaction." *Ergonomics.* **61**(10):1409–1427.
33. <sup>a</sup>Bigman YE, et al. (2019). "Holding Robots Responsible: The Elements of Machine Morality." *Trends Cogn Sci.* **23**(5):365–368.
34. <sup>a</sup>Saßmannshausen T, et al. (2021). "Trust in Artificial Intelligence Within Production Management - An Exploration of Antecedents." *Ergonomics.* **64**(10):1333–1350.
35. <sup>a</sup>Nass C, Moon Y (2000). "Machines and Mindlessness: Social Responses to Computers." *J Soc Issues.* **56**(1):81–103.
36. <sup>a</sup>Plebe A, Perconti P (2022). *The Future of the Artificial Mind.* Boca Raton: CRC Press.
37. <sup>a, b</sup>Safdari A (2025). "Toward an Empathy-Based Trust in Human-Otheroid Relations." *AI Soc.* **40**(5):3123–3138.

38. <sup>△</sup>Hassenzahl M, et al. (2020). "“Otherware” Needs Alternative Interaction Paradigms Beyond Naïve Anthropomorphism." *Nordic Conference on Human-Computer Interaction (NordiCHI)*. 25–29 October 2020, Tallinn, Estonia, accepted.
39. <sup>△</sup>Koivisto M, Grassini S (2023). "Best Humans Still Outperform Artificial Intelligence in a Creative Divergent Thinking Task." *Sci Rep*. **13**(1):13601.
40. <sup>△</sup><sup>a</sup>Cao B, et al. (2024). "On the Worst Prompt Performance of Large Language Models."
41. <sup>△</sup>Holzner N, Maier S, Feuerriegel S (2025). "Generative AI and Creativity: A Systematic Literature Review and Meta-Analysis."
42. <sup>△</sup>Meincke L, Mollick ER, Terwiesch C (2024). "Prompting Diverse Ideas: Increasing AI Idea Variance."
43. <sup>△</sup>Meincke L, et al. (2023). "Ideas Are Dimes a Dozen: Large Language Models for Idea Generation in Innovation." *SSRN Electron J*.
44. <sup>△</sup>Li JZ, et al. (2025). "Skill But Not Effort Drive GPT Overperformance Over Humans in Cognitive Reframing of Negative Scenarios." *PsyArXiv*.
45. <sup>△</sup>Salvi F, et al. (2025). "On the Conversational Persuasiveness of Large Language Models: A Randomized Controlled Trial." *Nat Hum Behav*. **9**(8):1645–1653. <http://arxiv.org/pdf/2403.14380v1>.
46. <sup>△</sup>Schoenegger P, et al. (2025). "Large Language Models Are More Persuasive Than Incentivized Human Persuaders."
47. <sup>△</sup>Hackenburg K, et al. (2025). "The Levers of Political Persuasion With Conversational Artificial Intelligence." *Science*. **390**(6777).
48. <sup>△</sup>Xu N, Ma X (2024). "LLM The Genius Paradox: A Linguistic and Math Expert's Struggle with Simple Word-Based Counting Problems."
49. <sup>△</sup>Yehudai G, et al. (2024). "When Can Transformers Count to N?"
50. <sup>△</sup>Glazer E, et al. (2024). "FrontierMath: A Benchmark for Evaluating Advanced Mathematical Reasoning in AI."
51. <sup>△</sup><sup>a</sup><sup>b</sup><sup>c</sup>Cheng L, Li X, Bing L (2023). "Is GPT-4 a Good Data Analyst?"
52. <sup>△</sup>Farquhar S, et al. (2024). "Detecting Hallucinations in Large Language Models Using Semantic Entropy." *Nature*. **630**(8017):625–630.
53. <sup>△</sup>Rathkopf C (2025). "Hallucination, Reliability, and the Role of Generative AI in Science."
54. <sup>△</sup>Poldrack RA, Lu T, Beguš G (2023). "AI-Assisted Coding: Experiments with GPT-4."
55. <sup>△</sup><sup>a</sup><sup>b</sup>Rae JW, et al. (2021). "Scaling Language Models: Methods, Analysis & Insights from Training Gopher."

56. <sup>a, b, c, d, e, f, g</sup>Bommasani R, et al. (2021). "On the Opportunities and Risks of Foundation Models."
57. <sup>Δ</sup>Agrawal A, McHale J, Oettl A (2023). "Superhuman Science: How Artificial Intelligence May Impact Innovation." *J Evol Econ.* 33(5):1473–1517.
58. <sup>Δ</sup>OpenAI, et al. (2023). "GPT-4 Technical Report."
59. <sup>Δ</sup>Touvron H, et al. (2023). "LLaMA: Open and Efficient Foundation Language Models."
60. <sup>Δ</sup>Yuan Z, et al. (2023). "How Well Do Large Language Models Perform in Arithmetic Tasks?"
61. <sup>a, b</sup>Nikankin Y, et al. (2024). "Arithmetic Without Algorithms: Language Models Solve Math With a Bag of Heuristics."
62. <sup>Δ</sup>He H (2025). "Defeating Nondeterminism in LLM Inference." Thinking Machines Lab.
63. <sup>a, b, c</sup>Bender EM, et al. (2021). "On the Dangers of Stochastic Parrots." In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. New York, NY, USA: ACM, 610–623.
64. <sup>a, b</sup>Bubeck S, et al. (2023). "Sparks of Artificial General Intelligence: Early Experiments With GPT-4."
65. <sup>a, b, c</sup>Doshi-Velez F, Kim B (2017). "Towards A Rigorous Science of Interpretable Machine Learning."
66. <sup>a, b</sup>Shinn N, et al. (2023). "Reflexion: Language Agents with Verbal Reinforcement Learning."
67. <sup>Δ</sup>Karimi P, et al. (2018). "Evaluating Creativity in Computational Co-Creative Systems."
68. <sup>Δ</sup>Weidinger L, et al. (2021). "Ethical and Social Risks of Harm from Language Models."
69. <sup>Δ</sup>Gil Y, Selman B (2019). "A 20-Year Community Roadmap for Artificial Intelligence Research in the US."
70. <sup>Δ</sup>Carlini N, et al. (2020). "Extracting Training Data From Large Language Models."
71. <sup>Δ</sup>Shumailov I, et al. (2023). "The Curse of Recursion: Training on Generated Data Makes Models Forget."
72. <sup>Δ</sup>Carolan K, Fennelly L, Smeaton AF (2024). "A Review of Multi-Modal Large Language and Vision Models."
73. <sup>Δ</sup>Merali A (2024). "Scaling Laws for Economic Productivity: Experimental Evidence in LLM-Assisted Translation."
74. <sup>Δ</sup>Sardana N, et al. (2024). "Beyond Chinchilla-Optimal: Accounting for Inference in Language Model Scaling Laws."
75. <sup>Δ</sup>Yauney G, Reif E, Mimno D (2023). "Data Similarity Is Not Enough to Explain Language Model Performance."
76. <sup>Δ</sup>Hou Z, et al. (2024). "Does RLHF Scale? Exploring the Impacts From Data, Model, and Method."
77. <sup>Δ</sup>Rafailov R, et al. (2023). "Direct Preference Optimization: Your Language Model Is Secretly a Reward Model."
78. <sup>Δ</sup>Xing S, et al. (2025). "LLMs Can Get 'Brain Rot'!"

79. <sup>a</sup><sub>b</sub>Liu NF, et al. (2023). "Lost in the Middle: How Language Models Use Long Contexts."
80. <sup>^</sup>Zhang L, et al. (2024). "SPRIG: Improving Large Language Model Performance by System Prompt Optimization."
81. <sup>a</sup><sub>b</sub>Sclar M, et al. (2023). "Quantifying Language Models' Sensitivity to Spurious Features in Prompt Design or: How I Learned to Start Worrying About Prompt Formatting."
82. <sup>a</sup><sub>b</sub>Zhuo J, et al. (2024). "ProSA: Assessing and Understanding the Prompt Sensitivity of LLMs."
83. <sup>^</sup>Fu Y, et al. (2023). "Specializing Smaller Language Models Towards Multi-Step Reasoning."
84. <sup>^</sup>Shojaee P, et al. (2025). "The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models Via the Lens of Problem Complexity."
85. <sup>^</sup>Lawsen A (2025). "Comment on The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models Via the Lens of Problem Complexity."
86. <sup>^</sup>Zhang Z, et al. (2023). "Unifying the Perspectives of NLP and Software Engineering: A Survey on Language Models for Code."
87. <sup>^</sup>Zheng Z, et al. (2023). "A Survey of Large Language Models for Code: Evolution, Benchmarking, and Future Trends."
88. <sup>^</sup>Ziabari AS, et al. (2025). "Reasoning on a Spectrum: Aligning LLMs to System 1 and System 2 Thinking."
89. <sup>^</sup>Zhang J, et al. (2025). "AdaptThink: Reasoning Models Can Learn When to Think."
90. <sup>^</sup>West P, et al. (2023). "The Generative AI Paradox: "What It Can Create, It May Not Understand"."
91. <sup>a</sup><sub>b</sub>Rajeev M, et al. (2025). "Cats Confuse Reasoning LLM: Query Agnostic Adversarial Triggers for Reasoning Models."
92. <sup>^</sup>Chatziveroglou G, Yun R, Kelleher M (2025). "Exploring LLM Reasoning Through Controlled Prompt Variations."
93. <sup>^</sup>Zheng Y, Zandsalimy M, Sushmita S (2025). "Behind the Mask: Benchmarking Camouflaged Jailbreaks in Large Language Models."
94. <sup>a</sup><sub>b</sub><sup>^</sup>Laban P, et al. (2025). "LLMs Get Lost In Multi-Turn Conversation."
95. <sup>^</sup>Abouzaid M, et al. (2026). "First Proof."
96. <sup>^</sup>Srivastava S, et al. (2024). "Functional Benchmarks for Robust Evaluation of Reasoning Performance, and the Reasoning Gap."
97. <sup>a</sup><sub>b</sub>Valmeekam K, et al. (2023). "On the Planning Abilities of Large Language Models: A Critical Investigation."



116. <sup>a</sup>Castelvecchi D (2024). "DeepMind Hits Milestone in Solving Maths Problems – AI's Next Grand Challenge." *Nature*. 632(8024):236–237.
117. <sup>a</sup>Majmudar N, Filatova E (2025). "Can LLMs Solve and Generate Linguistic Olympiad Puzzles?"
118. <sup>a</sup>Balunović M, et al. (2025). "MathConstruct: Challenging LLM Reasoning With Constructive Proofs."
119. <sup>a</sup>White C, et al. (2024). "LiveBench: A Challenging, Contamination-Limited LLM Benchmark."
120. <sup>a</sup>Jolicoeur-Martineau A (2025). "Less Is More: Recursive Reasoning With Tiny Networks."
121. <sup>a</sup>Fang YL, et al. (2025). "AI-Newton: A Concept-Driven Physical Law Discovery System Without Prior Physical Knowledge."
122. <sup>a</sup>Zhao A, et al. (2025). "Absolute Zero: Reinforced Self-Play Reasoning with Zero Data."
123. <sup>a</sup>Packer C, et al. (2023). "MemGPT: Towards LLMs as Operating Systems."
124. <sup>a, b</sup>Wu S, et al. (2024). "A Comparative Study on Reasoning Patterns of OpenAI's o1 Model."
125. <sup>a</sup>Behrouz A, Zhong P, Mirrokni V (2025). "Titans: Learning to Memorize at Test Time."
126. <sup>a</sup>Mitchell M (2021). "Why AI Is Harder Than We Think."
127. <sup>a, b, c, d, e</sup>Yang Q, et al. (2018). "Investigating How Experienced UX Designers Effectively Work with Machine Learning." In: Koskinen I, et al., eds. *Proceedings of the 2018 Designing Interactive Systems Conference*. New York, NY, USA: ACM, 585–596.
128. <sup>a, b, c, d, e</sup>Long D, Magerko B (2020). "What Is AI Literacy? Competencies and Design Considerations." In: Bernhaupt R, et al., eds. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, 1–16.
129. <sup>a</sup>Lin MP-C, et al. (2026). "Mapping AI Tools in Education: A Topic Modeling Analysis of Cognitive, Metacognitive, and Affective Insights." In: Graf S, Markos A, eds. *Generative Systems and Intelligent Tutoring Systems*. Cham: Springer Nature Switzerland, 88–101.
130. <sup>a, b</sup>Ruppert J, et al. (2024). "Taking Play and Tinkering Seriously in AI Education: Cases from Drag vs AI Teen Workshops." *Learn Media Technol.* 49(2):259–273.
131. <sup>a, b</sup>Reynolds L, McDonell K (2021). "Prompt Programming for Large Language Models: Beyond the Few-Shot Paradigm."
132. <sup>a, b, c, d</sup>Zamfirescu-Pereira JD, et al. (2023). "Why Johnny Can't Prompt: How Non-AI Experts Try (and Fail) to Design LLM Prompts." In: Schmidt A, et al., eds. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, 1–21.
133. <sup>a, b, c</sup>Wei J, et al. (2022). "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models."

134. <sup>△</sup>Errica F, et al. (2025). "What Did I Do Wrong? Quantifying LLMs' Sensitivity and Consistency to Prompt Engineering." In: Chiruzzo L, Ritter A, Wang L, eds. *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*. Stroudsburg, PA, USA: Association for Computational Linguistics, 1543–1558.
135. <sup>△</sup>Bsharat SM, Myrzakhan A, Shen Z (2023). "Principled Instructions Are All You Need for Questioning LLaMA-1/2, GPT-3.5/4."
136. <sup>△</sup>Schulhoff S, et al. (2024). "The Prompt Report: A Systematic Survey of Prompt Engineering Techniques."
137. <sup>△</sup>Basil S, et al. (2025). "Prompting Science Report 4: Playing Pretend: Expert Personas Don't Improve Factual Accuracy." [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=5879722](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5879722).
138. <sup>△</sup>Li C, et al. (2023). "Large Language Models Understand and Can Be Enhanced by Emotional Stimuli."
139. <sup>△</sup>Song I, et al. (2024). "ExploreSelf: Fostering User-Driven Exploration and Reflection on Personal Challenges with Adaptive Guidance by Large Language Models."
140. <sup>△</sup>Cheng X, Gao L, Luo X (2025). "From Emotion to Reflection: Leveraging EmotionPrompt Strategy to Empower Self-Determination in Decision-Making With Generative Artificial Intelligence." *Inf Manag.* 62(7):10419-4.
141. <sup>△</sup>Li L, et al. (2024). "Reflection-Bench: Evaluating Epistemic Agency in Large Language Models."
142. <sup>△</sup>Sahoo P, et al. (2024). "A Systematic Survey of Prompt Engineering in Large Language Models: Techniques and Applications."
143. <sup>△</sup>Meincke L, et al. (2025). "Prompting Science Report 2: The Decreasing Value of Chain of Thought in Prompting."
144. <sup>△</sup>Renze M, Guven E (2024). "Self-Reflection in LLM Agents: Effects on Problem-Solving Performance."
145. <sup>△</sup>Zhang W, et al. (2024). "Self-Contrast: Better Reflection Through Inconsistent Solving Perspectives."
146. <sup>△</sup>Xu S, et al. (2025). "Chain of Draft: Thinking Faster by Writing Less."
147. <sup>△</sup>Yin Z, et al. (2024). "Should We Respect LLMs? A Cross-Lingual Study on the Influence of Prompt Politeness on LLM Performance."
148. <sup>△</sup>Zhao S, Min Q (2025). "Can Polite Prompts Lead to Higher-Quality LLM Responses? - AI Theory of Mind Perspective." *Malaysia*. <https://aisel.aisnet.org/pacis2025/hci/hci/16>.
149. <sup>△</sup>Islam P, et al. (2023). "FinanceBench: A New Benchmark for Financial Question Answering."
150. <sup>△</sup>Suzuki Y, Banaei-Kashani F (2025). "Universe of Thoughts: Enabling Creative Reasoning with Large Language Models."

151. <sup>△</sup>Jahani E, et al. (2024). "Prompt Adaptation as a Dynamic Complement in Generative AI Systems."
152. <sup>△</sup>Acar OA (2024). "Beyond Prompt Engineering: Skills Marketers Need to Deploy Generative AI Successfully." *NIM Mark Intell Rev.* 16(1):18–23.
153. <sup>△</sup>Battle R, Gollapudi T (2024). "The Unreasonable Effectiveness of Eccentric Automatic Prompts."
154. <sup>△</sup>Zhang Q, et al. (2025). "Agentic Context Engineering: Evolving Contexts for Self-Improving Language Models."
155. <sup>△</sup>Kocielnik R, Amershi S, Bennett PN (2019). "Will You Accept an Imperfect AI?" In: Brewster S, et al., eds. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, 1–14.
156. <sup>△</sup>Jia J, et al. (2024). "Decision-Making Behavior Evaluation Framework for LLMs Under Uncertain Context."
157. <sup>△</sup>Rapp A, Di Lodovico C, Di Caro L (2025). "How Do People React to ChatGPT's Unpredictable Behavior? Anthropomorphism, Uncanniness, and Fear of AI: A Qualitative Study on Individuals' Perceptions and Understandings of LLMs' Nonsensical Hallucinations." *Int J Hum-Comput Stud.* 198:103471.
158. <sup>△</sup>Raisch S, Krakowski S (2021). "Artificial Intelligence and Management: The Automation–Augmentation Paradox." *Acad Manag Rev.* 46(1):192–210.
159. <sup>△</sup>Bubeck S, et al. (2025). "Early Science Acceleration Experiments With GPT-5."
160. <sup>△</sup>Mollick E (2024). *Co-Intelligence: The Definitive, Bestselling Guide to Living and Working with AI*. London: WH Allen.
161. <sup>△</sup>Maier S, Schneider M, Feuerriegel S (2025). "Partnering with Generative AI: Experimental Evaluation of Human-Led and Model-Led Interaction in Human-AI Co-Creation."
162. <sup>△</sup>Pareschi R (2024). "Beyond Human and Machine: An Architecture and Methodology Guideline for Centaurian Design." *Sci.* 6(4):71.
163. <sup>△</sup>Saghafian S, Idan L (2024). "Effective Generative AI: The Human-Algorithm Centaur."
164. <sup>△</sup>McCarthy IP, Hannigan T, Spicer A (2025). "The Risks of Botshit."
165. <sup>△</sup>Klein G (2001). *Sources of Power: How People Make Decisions*. 7th ed. Cambridge, Mass.: MIT Press.
166. <sup>△</sup>Klein G (2008). "Naturalistic Decision Making." *Hum Factors.* 50(3):456–460.
167. <sup>△</sup>Siu HC, et al. (2021). "Evaluation of Human-AI Teams for Learned and Rule-Based Agents in Hanabi."
168. <sup>△</sup>Jabarian B, Henkel L (2025). "Voice AI in Firms: A Natural Field Experiment on Automated Job Interviews."
169. <sup>△</sup>Wyszynski M, et al. (2025). "GenAI Improves Decision-Making: Revealing the Debiasing Capabilities of ChatGPT." *ICIS 2025 Proc.* 31. <https://aisel.aisnet.org/icis2025/gen-ai/gen-ai/31>.

170. <sup>△</sup>Chen Y, et al. (2023). "The Emergence of Economic Rationality of GPT." *Proc Natl Acad Sci U S A*. 120(51):e2316205120.
171. <sup>△</sup>Jiang L, et al. (2025). "Artificial Hivemind: The Open-Ended Homogeneity of Language Models (and Beyond)." *arXiv preprint arXiv:2501.08111*.
172. <sup>△</sup>Wenger E, Kenett Y (2025). "We're Different, We're the Same: Creative Homogeneity Across LLMs." *arXiv preprint arXiv:2501.08111*.
173. <sup>△</sup><sup>▭</sup>Meinke L, et al. (2025). "Call Me a Jerk: Persuading AI to Comply with Objectionable Requests." *Wharton Sch Res Pap*.
174. <sup>△</sup>Quattrociocchi W, Capraro V, Perc M (2025). "Epistemological Fault Lines Between Human and Artificial Intelligence." *arXiv preprint arXiv:2501.08111*.
175. <sup>△</sup>Fabiano F, et al. (2025). "Thinking Fast and Slow in Human and Machine Intelligence." *Commun ACM*. 68(8):72–79.
176. <sup>△</sup>Lockhart ENS (2025). "Creativity in the Age of AI: The Human Condition and the Limits of Machine Generation." *J Cult Cogn Sci*. 9(1):83–88.
177. <sup>△</sup>Brinkmann L, et al. (2023). "Machine Culture." *Nat Hum Behav*. 7(11):1855–1868.
178. <sup>△</sup>Ye Z, Li J (2024). "Artificial Intelligence Through the Lens of Metaphor: Analyzing the EU AIA." *Int J Digit Law Gov*. 1(2):361–381.
179. <sup>△</sup>Möck LA (2022). "Prediction Promises: Towards a Metaphorology of Artificial Intelligence." *J Aesthet Phenom*. 9(2):119–139.
180. <sup>△</sup>Gupta A, et al. (2024). "Assistant, Parrot, or Colonizing Loudspeaker? ChatGPT Metaphors for Developing Critical AI Literacies." *Open Praxis*. 16(1):37–53.
181. <sup>△</sup>Atef Y, Bali M (2025). "Whose Metaphor Is It Anyway? Analysing AI Metaphors From Positionality and Values of Speaker and Recipient." *Crit Stud Teach Learn*. 13(2).
182. <sup>△</sup><sup>▭</sup>Bory P, Natale S, Katzenbach C (2025). "Strong and Weak AI Narratives: An Analytical Framework." *AI & SOCIETY*. 40(4):2107–2117.
183. <sup>△</sup>Oldenburg N, Papyshv G (2025). "The Stories We Govern By: AI, Risk, and the Power of Imaginaries." *arXiv preprint arXiv:2501.08111*.
184. <sup>△</sup>Collins CR, Blood GW (1990). "Acknowledgment and Severity of Stuttering as Factors Influencing Nonstutterers' Perceptions of Stutterers." *J Speech Hear Disord*. 55(1):75–81.
185. <sup>△</sup>Boyle MP (2013). "Assessment of Stigma Associated With Stuttering: Development and Evaluation of the Self-Stigma of Stuttering Scale (4S)." *J Speech Lang Hear Res*. 56(5):1517–1529.

186. <sup>△</sup>Weizenbaum J (1966). "ELIZA—A Computer Program for the Study of Natural Language Communication Between Man and Machine." *Commun ACM*. **9**(1):36–45.
187. <sup>△</sup><sup>‡</sup>Kreps S, McCain RM, Brundage M (2022). "All the News That's Fit to Fabricate: AI-Generated Text as a Tool of Media Misinformation." *J Exp Polit Sci*. **9**(1):104–117.
188. <sup>△</sup>Fokas AS (2023). "Can Artificial Intelligence Reach Human Thought?" *PNAS Nexus*. **2**(12):pgad409.
189. <sup>△</sup>Sandini G, Sciutti A, Morasso P (2024). "Artificial Cognition vs. Artificial Intelligence for Next-Generation Autonomous Robotic Agents." *Front Comput Neurosci*. **18**:1349408.
190. <sup>△</sup>van Rooij I, et al. (2024). "Reclaiming AI as a Theoretical Tool for Cognitive Science." *Comput Brain Behav*. **7**(4):616–636.
191. <sup>△</sup>Wu C-K, et al. (2025). "Answer, Refuse, or Guess? Investigating Risk-Aware Decision Making in Language Models."
192. <sup>△</sup>Feng X, et al. (2024). "Large Language Model-Based Human-Agent Collaboration for Complex Task Solving."
193. <sup>△</sup>Chowa SS, et al. (2025). "From Language to Action: A Review of Large Language Models as Autonomous Agents and Tool Users."
194. <sup>△</sup><sup>‡</sup>Tankelevitch L, et al. (2023). "The Metacognitive Demands and Opportunities of Generative AI."
195. <sup>△</sup>Kahneman D (2013). *Thinking, Fast and Slow*. New York: Farrar Straus and Giroux.
196. <sup>△</sup><sup>‡</sup><sup>§</sup>Shin D (2025). *Debiasing AI*. New York: Routledge.
197. <sup>△</sup>Gerlich M (2025). "AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking." *Societies*. **15**(1):6.
198. <sup>△</sup>Guinrich RE, Mehta D, Bhatt U (2026). "Belief Offloading in Human-AI Interaction."
199. <sup>△</sup>Shen JH, Tamkin A (2026). "How AI Impacts Skill Formation."
200. <sup>△</sup>Fernandes D, et al. (2026). "AI Makes You Smarter But None the Wiser: The Disconnect Between Performance and Metacognition." *Comput Hum Behav*. **175**:108779.
201. <sup>△</sup>Jarrahi MH (2018). "Artificial Intelligence and the Future of Work: Human-AI Symbiosis in Organizational Decision Making." *Bus Horiz*. **61**(4):577–586.
202. <sup>△</sup>Sadeghian S, et al. (2025). "WorkAI: A Toolkit for the Design of AI-Driven Future of Work." *Proc ACM Hum-Comput Interact*. **9**(7):1–27.
203. <sup>△</sup>Davis GF (2019). "How to Communicate Large-Scale Social Challenges: The Problem of the Disappearing American Corporation." *Proc Natl Acad Sci U S A*. **116**(16):7698–7702.

204. <sup>△</sup>Burggräf P, Wagner J, Saßmannshausen TM (2021). "Sustainable Interaction of Human and Artificial Intelligence in Cyber Production Management Systems." In: B.-A. Behrens, et al., eds. *Production at the Leading Edge of Technology*. Berlin, Heidelberg: Springer Berlin Heidelberg, 508–517.
205. <sup>△</sup>van Quaakebeke N, Gerpott FH (2023). "The Now, New, and Next of Digital Leadership: How Artificial Intelligence (AI) Will Take Over and Change Leadership as We Know It." *J Leadersh Organ Stud*. 30(3):265–275.
206. <sup>△</sup>Cvetkovic A, et al. (2024). "Do We Trust Artificially Intelligent Assistants at Work? An Experimental Study." *Hum Behav Emerg Technol*. 2024:1–12.

## Declarations

**Funding:** No specific funding was received for this work.

**Potential competing interests:** No potential competing interests to declare.