Conceptualizing Agency: A Framework for Human-AI Interaction

Alexander Zhang[1], Lav R. Varshney[2]

[1]University of Illinois Urbana-Champaign, Department of Philosophy, USA
[2]University of Illinois Urbana-Champaign, Department of Electrical and Computer Engineering, USA

Abstract

This paper proposes a framework for understanding *intentional agency* in the context of increasingly autonomous AI systems. As AI tools perform complex tasks independently, they challenge our traditional notions of agency—particularly the ability to act intentionally, grounded in what Anscombe calls "practical knowledge." We explore intentional agency through three lenses: (1) Anscombe's concept of intention, where we examine how AI tools complicate human intentionality and our ability to give reasons for actions; (2) shared intention, which frames agency in joint human activities; and (3) relational intention, a novel synthesis drawing on Anscombe and Fiske's Relational Models Theory which captures how relationships shift between different relational models. Our analysis reveals that agency is not merely an individual capacity but is fundamentally shaped by relationships with others who can interpret and respond to our reasons. AI systems, however, disrupt this dynamic: while they assist in tasks, they lack the intersubjectivity required to engage in mutual reason-giving, potentially diminishing our practical knowledge and intentional agency. We conclude that AI should be designed to support human-human relationships rather than supplant them, with humans retaining relational authority in human-AI interactions to preserve agency.

Keywords

Human-AI co-creation, Agency, Generative AI, Relational Models Theory, Anscombe

1. Introduction

Agency is fundamental to human-computer interaction (HCI) and moral philosophy alike. Typically defined as the capacity to act intentionally, it lies at the heart of human experience, yet its details remain elusive, particularly as we integrate advanced AI systems into our lives. HCI researchers strive to safeguard it amid technological advances, but what exactly are we protecting? This question grows urgent with the rise of AI tools that perform tasks autonomously, often beyond our immediate oversight. How should we rethink agency in this landscape? And how do these tools affect it differently than human interactions?

Notions of agency and autonomy have long been central to HCI research, particularly in creative contexts, where intentionality is pivotal to understanding creativity [1, 2]. Recent work also stresses the need for intentionality in creativity, suggesting that intentionality is often essential if we are to consider computer outputs genuinely creative [3]. Yet traditional frameworks for understanding agency may not fully capture the complexities of human-AI

Joint Proceedings of the ACM IUI Workshops 2025, March 24-27, 2025, Cagliari, Italy
azhan7@illinois.edu (A. Zhang); varshney@illinois.edu (L. R. Varshney)

interaction.

Traditionally, agency has been framed as the ability to act, what Dai [4] terms "mechanistic agency." While useful, this view is narrow, focusing on outcomes rather than the reasons behind them. Dai's alternative, "volitional agency," centers on intrinsic motivations, but we propose a different perspective, inspired by G.E.M. Anscombe's work on intention [5]. Anscombe highlights "practical knowledge"—knowing why we act without observing ourselves—as central to intentional action. More than just performing acts, agency involves understanding our actions in ways explicable to ourselves and others. We argue that this explanatory process is inherently relational, grounded in our connections with those who share a context of reason-giving.

Building on Anscombe, we suggest that intentional agency emerges not in isolation but through relationships. Fiske's Relational Models Theory (RMT) [6] provides a lens to see how these relations—whether authority-based, communal, or otherwise—shape our capacity to act intentionally. AI complicates this picture. By assisting in tasks, it can enhance our mechanistic agency (e.g., boosting productivity), much like human interactions do. But unlike humans, AI lacks the intersubjectivity—the mutual recognition of consciousness—that underpins intentional agency. Without an authentic "other" to engage in reason-giving, AI risks isolating us, eroding our practical knowledge, and thus reducing our agency.

This paper develops a framework for understanding intentional agency through three dimensions of intention. First, we discuss Anscombe's notion of intention, showing how AI's autonomy and opacity challenge practical knowledge. Next, we explore shared intention, examining how human joint activities depend on mutual recognition between authentic "others"—something AI cannot provide as a mere tool. Finally, we introduce "relational intention," arguing that agency hinges on dynamic, intersubjective relationships—something AI cannot fully replicate. We aim to deepen the philosophical understanding of agency in an AI-driven world and to guide the design of systems that bolster, rather than undermine, human relationships and intentional action.

2. Intention: A First Sketch

2.1. Anscombe's *Intention*

A man sketches a pencil drawing of a robot. His hand moves deliberately, sketching the curved lines of a robot's smiling face. His movement is no reflex—like breathing or blinking—but an *intentional* action. What makes it so? And what is his intention? G.E.M. Anscombe, in her seminal work *Intention*, offers a way to understand this through a chain of "why" questions: "Why are you moving your pencil across the page?" "To draw a robot." "Why draw a robot?" "To express my enthusiasm about human-AI interaction" "Why?" The questioning continues, connecting actions to their deeper goals until we reach the intention with which he acts—perhaps to visually celebrate the creative potential of human-AI interaction. An intentional action, then, is an action where the question "Why?" finds application. Anscombe notes that there can be motivations beyond this, but that these are not necessarily descriptions of the agent's actions. For example, consider that, secretly, the man believes that in the future, robots will dominate humans, and that this artistic appreciation might curry favor for him. But such forward-looking explanations exceed the immediate intention with which he acts.

Critically, for Anscombe, intentional action hinges on what she calls "practical knowledge": knowing what we're doing and why, without needing to observe ourselves as we might a stranger's behavior. Schwenkler [7] summarizes this into three criteria, building on Anscombe:

1. A person acts intentionally only if they know they are doing it.
2. They know this without observing their actions.
3. They don't need observation to explain why they're acting.

The sketcher knows he's moving his pencil across the page to draw a robot not because he watches his hand, but because he grasps his purpose directly.

2.2. AI Explainability and the Challenge to Practical Knowledge

Now imagine he uses an AI tool to generate parts of the sketch—say, the robot's arm. He prompts the system, and it produces an effect. Asked "Why did you shade it that way?" he might pause: "I don't know—that's just how the AI rendered it." Unlike watercolor's unpredictable bleed, which he could still explain through technique, the AI's output feels opaque. It performs the task autonomously, but its process eludes his practical knowledge. By failing to know what he's doing and why without observation, his use of AI to perform a task results in a reduction of practical knowledge. Could we design AI to be more accessible to our practical knowledge?

Selbst and Barocas [8] identify two hurdles here: inscrutability and nonintuitiveness. Inscrutability points to the technical opacity of AI and its complexity that defies human comprehension. More pressing for Anscombe's framework is nonintuitiveness: even if the system's workings were transparent, its outputs might not align with reasons humans find meaningful. An AI might shade a drawing in a mathematically optimal way, but if the sketcher can't connect that to his intent, it fails to inform his "why."

This distinction between technical transparency and human-meaningful reasons reveals something crucial about intentional action. When we ask "why" questions about human actions, we expect answers that reference goals, motivations, and broader contexts, not detailed descriptions of neural activity or muscle movements. Similarly, while we might understand precisely how an AI model processes data (addressing inscrutability), this differs from understanding why it took certain actions in terms that connect to human intentions (addressing nonintuitiveness).

Moreover, the ability to give reasons for our actions develops through social contexts where meanings are collectively constructed and validated among conscious beings. Our sketcher can explain his shading technique because he learned its significance within a community of artists who share understanding about how certain techniques evoke certain effects. AI systems fundamentally cannot participate in these meaning-making social contexts. Even when they provide convincing explanations, these remain simulations of reason-giving rather than genuine participation in the shared world of human meaning.

This points to a deeper problem than mere technical opacity. In traditional tool use, intention flows clearly from the human actor—we wouldn't ask a pencil "why" it moved a certain way. But AI introduces autonomous decision-making that distributes agency in ways that can obscure our own practical knowledge. When the AI generates unexpected variations or makes autonomous choices, it doesn't just assist our action but potentially disconnects us from direct understanding of what we're doing. Unlike a brush stroke we can explain through technique, or a collaborator's

choice we can understand through shared intention, the AI's contributions remain fundamentally opaque to the kind of reason-giving that characterizes intentional action.

2.3. The Role of the "Other"

This opacity points to something fundamental about intentional action itself: Anscombe's "why" questions aren't solitary reflections; rather, they imply an "other" who can interpret and challenge our reasons within a shared world of meaning. Practical knowledge is made meaningful by being publicly intelligible, shaped by a context where reasons are given and received. The sketcher doesn't just know he's moving his pencil to draw a robot—he can also explain his actions to a teacher or friend who understands the craft.

Intentional agency, then, isn't merely self-directed; it presupposes an "other" who anchors our reasons in a human web of understanding. AI tools, however sophisticated, can't fill this role. They generate outputs—sometimes stunning ones—but lack the consciousness to engage in mutual reason-giving. When the sketcher admits, "That's just what the AI did," there's no "other" to respond, refine, or contest his intent. The tool enhances his output but isolates him from the relational exchange that sustains agency. This absence doesn't just obscure practical knowledge; it threatens the intersubjective ground of intentional action itself—a theme we'll expand in the next section.

3. Shared Intention: Collaborating with the "Other"

3.1. Practical Knowledge in Shared Agency

A playwright and an actor are creating a play together. They're not just doing distinct, separate actions, described with different intentions; they're sharing the intention of "creating a play." But how do we understand this transition from individual intentions to shared intentions? Traditional accounts of shared agency often struggle with this question, sometimes appealing to problematic notions of a "group mind"—a collective consciousness somehow emerging from or hovering above individual actors. Such approaches raise metaphysical issues: What is the relationship between individual and group consciousness? How could such a collective mind arise? These questions become especially troublesome when considering artificial "agents."

Laurence [9], drawing on Anscombe, offers a more grounded approach. Rather than positing mysterious collective mental states, he explains shared agency through patterns of knowledge and reason-giving. In collective action, different participants may have different levels of understanding about what they're doing together, but the activity must meet three key criteria:

- Some suitably placed member(s) must know without observation what purposes the group is pursuing. This sustains the action's intentionality.
- Other members may epistemically defer to these leaders while still participating in genuine collective action. They might not grasp the full plan but understand their actions in relation to the group's purpose.
- All members' actions must be explicable through reference to the group's activity. Each contribution can be rationalized by the collective aim.

This framework illuminates the shared intention between the playwright and the actor. For each of them, the "other" isn't some distant observer but rather a collaborator. Each collaborator knows without observation what the pair is aiming for: putting on a play. The actor might not grasp every aspect of the script—she's focused on her cues—but she *epistemically defers* to the playwright, trusting her acting aligns with the whole. Their actions (writing, acting) cohere because they serve the play. The playwright's script provides the actor's lines; her delivery informs the playwright's revisions.

"Epistemic deferment" is a crucial term that Laurence uses to understand shared agency. It occurs in situations involving authority or trust, where an individual agent does not know without observation the explanation of their action, and "the epistemic buck will have to be passed to some other suitably placed member of the group who does" [9]. While collaborating may reduce certain aspects of one's practical knowledge, this is compensated for through epistemic deferment—allowing us to rely on another's knowledge of shared intention while retaining our agency. This capacity for epistemic deferment is what enables genuinely shared agency in the first place.

Consider another, larger example of how epistemic deferment plays out during shared agency on a ship's crew during a storm: the captain knows the goal—reaching safe harbor—without observing every rope; the first mate translates orders into tasks; sailors haul lines, explaining their efforts by reference to "what we're doing" as a ship, even if they don't see the entire plan. Some defer to others' knowledge, binding individual acts into collective action. The unity emerges not from shared consciousness but from each member's ability to situate their actions within the group's purpose through genuine epistemic deferment.

3.2. The Illusion of an "Other"

Now suppose the playwright replaces the actor with an AI tool. He feeds it a script, and it generates a performance—say, a voice reciting lines with pauses and inflection. Does this constitute shared agency? The tool produces a reading, but it's not deferring to the playwright or apprehending a joint aim—it's merely outputting what was prompted, or what its training yields. The playwright's practical knowledge persists only if he can explain its contribution, as he could for the actor. If not, the "shared" aspect falters. Unlike the crew, where deferment unites us, the AI provides no such connection.

Rödl [10] articulates this precisely: joint action requires "consciousness of consciousness"—a mutual awareness where we each understand our own part and the other's comprehension of the shared endeavor. This goes beyond mere "theory of mind" or the ability to model others' mental states. What's required is genuine intersubjectivity: not just predicting what others might think or do, but participating in a shared world of meaning where each recognizes the other as a fellow subject. Returning to the play: the playwright knows he's writing this scene to set the mood; the actor knows she's acting it to convey the rhythm; and they both know the other knows this constructs the play. Rödl terms this "plural self-consciousness"—a shared awareness uniting our efforts into one intention.

An AI tool reciting my script falls short. It can produce a performance—perhaps an effective one—based on my words or a data set, but there's no mutual recognition, no reciprocal exchange where it grasps my intent and the user grasp its. While it might simulate understanding through

sophisticated response patterns, this differs fundamentally from the intersubjective awareness that characterizes human collaboration. It lacks consciousness, let alone consciousness of mine.

This isn't a matter of refining AI to simulate that "otherness" better. Even if an AI could perfectly model human mental states, this would remain different from genuine intersubjectivity. Enhancing its simulation only isolates one further, transforming a joint effort into a solitary task with aids. Human shared intention adapts through epistemic deferment; AI-assisted tasks remain static, leaving us alone with the result and no "other" to defer to.

Thus, AI's effect on agency extends beyond its impact on practical knowledge, which human collaborators can also influence. What fundamentally distinguishes AI is that while we can meaningfully defer to human partners' knowledge and rely on their participation in shared intentions, AI cannot authentically fill this role due to its inherent lack of consciousness. The next section examines how epistemic deferment manifests across different relational models, revealing both how human relationships constitute agency and the concerning implications of attempting to epistemically defer to an artificial system incapable of genuine reciprocity.

4. Relational Intention: Dancing with the "Other"

Rödl writes, "The current literature on joint action for the most part concerns itself with action by equals. This may be right; perhaps this is the fundamental case of joint action." While Rödl highlights the traditional focus on equal partners in joint action [10], we must recognize that human relationships transcend simple categorization. They are better illustrated as dance—fluid, dynamic, and inherently relational.

Dance embodies this complexity as a "distinctively human force of intelligence," where meaning emerges through movement and interaction [11]. As philosopher Alva Noë observes: "We dance from the moment we're born to the moment we die... Everything we do is structured in the way a dance is structured—a dynamic of interaction, a conversation, a relationship. In the large and small. A society. There's just a constant distributed give and take and flow" [12].

This metaphor of dance describes not just intelligence but also agency. We never rest in fixed relational models with others but constantly shift and adapt. Consider how relationships evolve: a child growing with their parents, friends developing trust, colleagues learning to collaborate. Each interaction involves risk, adjustment, and mutual growth. Partners naturally shift roles—leading in one moment, following in another. Even in formal structures, like a ship's crew, relationships remain dynamic. The captain might issue direct orders during a storm, engage in tactical planning with the first mate, or welcome creative problem-solving from experienced sailors. This fluidity works because all parties remain conscious subjects who can genuinely engage in reason-giving and mutual recognition.

While recognizing it may be hard to capture precisely the dynamics of the inherent fluidity of human relationships, Fiske's Relational Models Theory [6]helps to gain an understanding of what different relationship dynamics are. He identifies four primary models through which relationships develop, although we often use combinations of these models in our relationships:

- **Authority Ranking:** Hierarchical relationships where authority and deference flow clearly
- **Market Pricing:** Exchange-based relationships focused on mutual benefit

- **Communal Sharing:** Close bonds characterized by shared identity and resources
- **Equality Matching:** Balanced reciprocity where contributions and benefits are matched

Through these relational models, we discover distinct patterns of intention and epistemic deferment. These patterns shape how we explain and understand actions in relation to others. In **Authority Ranking**, practical knowledge flows hierarchically through epistemic deferment, where crew members may not grasp the full strategic picture but maintain intentional agency by deferring to the captain's knowledge of the mission. In **Market Pricing**, intentions interlace through mutual epistemic deferment based on expertise, as partners maintain agency by each deferring to the other's knowledge while preserving their own practical understanding. In **Communal Sharing**, epistemic deferment occurs naturally through deep familiarity, where family members fluidly share practical knowledge, each trusting the others' understanding of shared purposes. In **Equality Matching**, intentions align through balanced epistemic deferment, as friends maintain individual agency while trusting each other's judgment equally, creating reciprocal patterns of knowledge-sharing.

4.1. The Limits of Simulated Relations

However, as AI systems become more sophisticated, they challenge this understanding of relational dynamics. AI fundamentally disrupts relational patterns in two critical ways. First, it can diminish our practical knowledge without providing the authentic "other" to whom we could meaningfully defer. Second, when perceived as too authoritative, it can reverse the intended dynamic, subtly eroding human dignity and agency - what Malone et al. describe as an "epistemic harm" that fails to recognize human workers' expertise and constrains their ability to act independently [13].

Empirical work by Tschopp et al. [14] suggests that while humans attempt to engage with AI through these same relational models, Authority Ranking—a hierarchical dynamic—often dominates these interactions. This raises crucial questions about agency. Their research combines Fiske's Equality Matching and Communal Sharing models into a single "Peer Bonding" category when examining human-AI interaction. Each model of interaction breaks down distinctly with AI. In **Authority Ranking**, AI appears to defer to human commands but cannot genuinely grasp or adapt to our purposes. In **Market Pricing**, AI simulates exchange but cannot truly assess or negotiate value. In **Peer Bonding**, AI most dangerously mimics collaboration, generating responses that feel creative and engaged.

Yet regardless of relational model, AI cannot provide the foundation for genuine epistemic deferment. When one works with a human collaborator, their shared understanding allows them to act intentionally even when one doesn't fully grasp every aspect of their joint action. One can defer to the other's judgment because they can genuinely engage in reason-giving and mutual recognition. With AI, this foundation is missing. It can generate impressive outputs, but we cannot meaningfully defer to its "understanding" because it has none—it's an artifact, not an "other."

This reveals a paradox: AI tools designed to enhance human capability can inadvertently diminish human agency by disrupting the relational foundation of intentional action. When we treat AI as an authentic "other" or allow it to assume too much authority, we risk losing not just

practical knowledge but the dignity that comes from being recognized as an intentional agent by fellow conscious beings.

Considering Lubart's [15] typology of AI roles in creative processes, our analysis suggests that AI should serve as a facilitator of human communication and relationships rather than a simulated creative partner. While AI can act as nurturer, coach, colleague or penpal, the latter role—enhancing connections between human creators—best preserves intentional agency. When AI assumes the role of colleague, it risks eroding the authentic "otherness" that grounds human creative collaboration, potentially diminishing rather than enhancing creative agency. Instead, AI should be designed to strengthen the bonds between human creators while remaining clearly positioned as a tool rather than a relational partner.

4.2. Design Principles

Drawing from our theoretical analysis of intentional agency and relational models, we propose several preliminary design considerations that warrant empirical investigation:

- **Preserve Human Relationality:** Systems should be designed to strengthen human-human bonds rather than replace them. For instance, AI writing assistants might facilitate collaborative editing between authors rather than generating content in isolation. As Gaudet et al. [16] warns, "we should safeguard against AI being used, intentionally or unintentionally, as a substitute for genuine relationships. AI should only be used to enrich our relationships with one another, our community, and with nature." Further, "there should always be an easy route to appeal an AI decision to a person" [16].
- **Maintain Clear Authority Structures:** While AI can assume various relational models, systems should be designed to preserve human agency and prevent unintended authority reversals. This might involve explicit user control over when and how AI assistance is activated.
- **Support Practical Knowledge:** Interfaces should help users maintain understanding of their actions when using AI tools, potentially through explanatory features or deliberate moments of reflection. Perhaps the surest way for users of AI to preserve practical knowledge when doing a task is to "initially learn how to do all the steps of a process before they automate it with AI. Then they can become aware of what they are outsourcing, and where the AI could err" [16].
- **Calibrate Trust Appropriately:** Systems should explicitly communicate their limitations and uncertainty. In Market Pricing interactions, this means quantifying uncertainty in predictions. In Peer Bonding scenarios, systems should acknowledge their inability to truly engage in creative collaboration, helping maintain appropriate expectations.

For individual users of AI, we suggest reflecting on:

- Could this task benefit from human collaboration instead of AI assistance?
- What relational model am I defaulting to with the AI, and is it appropriate?
- How is this tool affecting my ability to explain and understand my own actions?
- Am I calibrating my trust appropriately, or treating AI simulation as genuine relationship?

These proposals stem from our theoretical framework but require empirical validation. Future work should systematically test how different design approaches affect human agency, relationship quality, and practical knowledge while maintaining appropriate levels of trust across different models of interaction.

5. Future Work

This framework opens several promising avenues for future research, both philosophical and empirical.

Philosophical Investigations:

- **Foundations of Intentional Agency:** While we've argued that intentional action implies an "other," this needs deeper examination. Can genuine intentional action occur in isolation, or does it always presuppose a relational context?
- **The Nature of "Otherness":** We need stronger theoretical grounds for why AI can or cannot constitute an authentic "other." This connects to fundamental questions in philosophy of mind and intersubjectivity, as well as the tricky problem of AI 'consciousness.'

Empirical Research:

- **Measuring Practical Knowledge:** Develop methodologies to assess how different AI interactions affect users' ability to explain their actions without observation.
- **Relational Dynamics:** Study how different interface designs influence perceived relational models and their impact on human agency.
- **Collaborative Systems:** Test designs that prioritize human-human relationships while incorporating AI assistance.

Conclusion

As AI systems grow more sophisticated, they challenge our understanding of agency, intention, and human relationship. Our analysis reveals that agency isn't merely an individual capacity but emerges through authentic encounters with others who can recognize and respond to our intentions. This "otherness" is what makes genuine human connection possible—and what AI, despite its capabilities, cannot replicate. The danger lies in mistaking AI's simulated relationships for authentic ones. When we read a text, we encounter its author's consciousness; when we use AI, we interact with a system that can process but not understand. The temptation to replace difficult human relationships with frictionless AI interactions threatens the very foundation of intentional agency.

Through our investigation of individual, shared, and relational intention, we've seen how agency emerges from and depends upon authentic human relationships. Whether in direct collaboration or more complex social dynamics, intentional action requires genuine connection with conscious others who can understand and respond to our reasons.

As we go forwards in an AI-filled future, we must preserve and strengthen these human bonds rather than replace them with convenient but hollow substitutes. The hope of respecting

human agency lies not in perfecting AI's simulation of "otherness" but in using it to support the authentic human relationships that make intentional action possible.

Acknowledgments

We thank John Schwenkler, Kush Varshney, and the reviewers of the HAI-GEN workshop for their invaluable feedback.

References

[1] D. Bennett, O. Metatla, A. Roudaut, E. D. Mekler, How does hci understand human agency and autonomy?, in: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, CHI '23, ACM, 2023, pp. 1–18. URL: http://dx.doi.org/10.1145/3544548.3580651. doi:10.1145/3544548.3580651.

[2] T. Lubart, How can computers be partners in the creative process: Classification and commentary on the special issue, International Journal of Human-Computer Studies 63 (2005) 365–369. doi:10.1016/j.ijhcs.2005.04.002.

[3] L. R. Varshney, Limits theorems for creativity with intentionality, in: Proceedings of the 11th International Conference on Computational Creativity (ICCC'20), 2020, pp. 390–393.

[4] J. Dai, Beyond personhood: Agency, accountability, and the limits of anthropomorphic ethical analysis, 2024. URL: https://arxiv.org/abs/2404.13861. arXiv:2404.13861.

[5] G. E. M. Anscombe, Intention, Harvard University Press, Cambridge, 1957.

[6] A. P. Fiske, The four elementary forms of sociality: Framework for a unified theory of social relations, Psychological Review 99 (1992) 689–723. doi:10.1037/0033-295X.99.4.689.

[7] J. Schwenkler, Anscombe's *Intention*: A Guide, Oxford University Press, New York, USA, 2019.

[8] A. D. Selbst, S. Barocas, The intuitive appeal of explainable machines, Fordham Law Review 87 (2018-2019) 1085.

[9] B. Laurence, An anscombian approach to collective action, in: A. Ford, J. Hornsby, F. Stoutland (Eds.), Essays on Anscombe's Intention, Harvard University Press, 2011.

[10] S. Rödl, Joint action and plural self-consciousness, Journal of Social Philosophy 49 (2018) 124–136. doi:10.1111/josp.12226.

[11] A. Noë, The entanglement, Harper's Magazine (2023). URL: https://harpers.org/archive/2023/07/alva-noe-the-entanglement/, accessed: 2025-02-25.

[12] Alva Noë, Dance As A Way Of Knowing: Interview With Alva Noë, 2012. URL: https://www.youtube.com/watch?v=FbWVERm5bsM, interview by dancetechtv, Published on YouTube, July 25, 2012.

[13] E. Malone, S. Afroogh, J. DCruz, K. R. Varshney, When trust is zero sum: Automation threat to epistemic agency, 2024. URL: https://arxiv.org/abs/2408.08846. arXiv:2408.08846.

[14] M. Tschopp, M. Gieselmann, K. Sassenberg, Servant by default? how humans perceive their relationship with conversational ai, Cyberpsychology: Journal of Psychosocial Research on Cyberspace 17 (2023) Article 9. URL: https://cyberpsychology.eu/article/view/21003. doi:10.5817/CP2023-3-9.

[15] T. Lubart, How can computers be partners in the creative process: classification and commentary on the special issue, Int. J. Hum.-Comput. Stud. 63 (2005) 365–369. URL: https://doi.org/10.1016/j.ijhcs.2005.04.002. doi:10.1016/j.ijhcs.2005.04.002.

[16] M. J. Gaudet, P. Scherz, N. Herzfeld, J. J. Wales, N. Colaner, J. Coogan, M. Courtois, B. Cutter, D. E. DeCosse, J. C. Gable, B. Green, J. Kintz, C. A. Labrecque, C. Moon, A. Ramelow, J. P. Slattery, A. M. Vega, L. G. Vera, A. Vicini, W. von Eschenbach, Encountering Artificial Intelligence: Ethical and Anthropological Reflections, Pickwick Press, Eugene, OR, 2023.