

Language Models as Critical Thinking Tools: A Case Study of Philosophers

Andre Ye^{αγ}, Jared Moore^β, Rose Novick^γ, Amy X. Zhang^α

andrey@uw.edu, jlcmoore@stanford.edu, amnovick@uw.edu, axz@cs.washington.edu

^α Paul G. Allen School of Computer Science and Engineering, University of Washington

^β Department of Computer Science, Stanford University

^γ Department of Philosophy, University of Washington

Abstract

Current work in language models (LMs) helps us speed up or even skip thinking by accelerating and automating cognitive work. But can LMs help us with *critical thinking* — thinking in deeper, more reflective ways which challenge assumptions, clarify ideas, and engineer new concepts? We treat philosophy as a case study in critical thinking, and interview 21 professional philosophers about how they engage in critical thinking and on their experiences with LMs. We find that philosophers do not find LMs to be useful because they lack a sense of selfhood (memory, beliefs, consistency) and initiative (curiosity, proactivity). We propose the *selfhood-initiative* model for critical thinking tools to characterize this gap. Using the model, we formulate three roles LMs could play as critical thinking tools: the Interlocutor, the Monitor, and the Respondent. We hope that our work inspires LM researchers to further develop LMs as critical thinking tools and philosophers, and other ‘critical thinkers’ to imagine intellectually substantive uses of LMs.

1 Introduction

“But I like the inconveniences.” — “We don’t,” responds the Controller. “We prefer to do things comfortably.” — “But I don’t want comfort,” John gasps. “I want God, I want poetry, I want real danger, I want freedom, I want goodness. I want sin.” — “In fact,” says the Controller, “you’re claiming the right to be unhappy ... the right to live in constant apprehension of what may happen tomorrow; ... the right to be tortured by unspeakable pains of every kind.” There is a long silence. “I claim them all,” says John at last. (*Minimally adapted from Huxley (1932).*)

Language Models (LMs) have recently alleviated a whole host of our intellectual inconveniences. They can help us do things we would have begrudgingly done by ourselves otherwise: write code (Chen et al., 2021; Rozière et al., 2023), generate emails (Goodman et al., 2022), and translate text (Costa-jussà et al., 2022). In sparking ideas by generating stories (Schwitzgebel et al., 2023) and concept designs (Cai et al., 2023), LMs offer shortcuts to gaining new thoughts. They also help us put our thinking into words by revising (Mysore et al., 2023) and giving feedback (Liang et al., 2024) on our writing. In all these cases, LMs help us speed up and circumvent the inconveniences of thinking ourselves.

In many contexts, however, the “inconvenience” of thinking is not a temporary problem to be alleviated, but a deep puzzle to be reflected upon. Many people are invested in specific areas of intellectual inquiry — e.g., historians, scientists, philosophers — and more generally, in reflection and engagement with the world — e.g., as informed political citizens, critical information consumers, and moral actors. They are interested in identifying and challenging assumptions, clarifying muddled ideas, and engineering new and useful ways to think. Core to this sort of inquiry is *critical thinking* — “the propensity and skill to engage in an activity with reflective skepticism” (McPeck, 2016). Can LMs serve as tools for *critical*

thinking — helping us think more deeply and in more complex ways, rather than faster or not at all? What if — like John — *we claim all the rights to think* (Bućinca et al., 2021)?

To investigate how LMs can serve as critical thinking tools, we use philosophers as a case study — philosophers being people who are in the business of thinking critically about a wide range of concepts and ideas. We interview **21 professional philosophers** to understand their thinking processes, collect their experiences with and views on current LMs, and brainstorm the roles LMs could play as critical thinking tools in philosophy (§3). We find that current philosophers *do not* think LMs are good critical thinking tools (§4) for two primary reasons: LMs are too neutral, detached, and nonjudgmental (§4.2); and LMs are too servile, passive, and incurious (§4.3). We propose the *selfhood-initiative* model for critical thinking tools, which explains why philosophers find conversations with other philosophers and reading philosophical texts to be more helpful for their work than current LMs (§5.1). Using the model, we describe **three roles** LMs could play as critical thinking tools: the Interlocutor, the Monitor, and the Respondent (§5.2). Finally, we outline how these LMs could inform metaphilosophical questions and shape the discipline of philosophy (§6.3), and discuss challenges in building LMs (§6.1) and interfaces (§6.2) for critical thinking.

2 Background and Related Work

2.1 Language Models as Thinking Tools

A large and growing literature investigates how LMs can serve as thinking tools for humans engaged in intellectual work. This research tends to concern how LMs can serve two intellectual functions: *idea stimulation* (roughly, “divergent thinking”) and *idea refinement* (roughly, “convergent thinking”) (Banathy, 1996; Design Council, 2019).

LMs can provide *stimulus for ideas* — information and (re)formulations which provoke and guide creative processes. Generally, LMs can expand idea sets (Fede et al., 2022), produce creative analogies (Bhavya et al., 2023) and metaphors (Chakrabarty et al., 2021); discover concepts (Lam et al., 2024), and facilitate group ideation (Rayan et al., 2024; Shaer et al., 2024). LMs to open-endedly propose plots, characters, and entire stories for creative writers (Calderwood et al., 2020; Schmitt & Buschek, 2021; Yuan et al., 2022; Mirowski et al., 2023; Chung et al., 2022; Chakrabarty et al., 2023); but also provide inspiration in more constrained creativity tasks, such as science writing (Gero et al., 2021; Kim et al., 2023). Although fraught with pitfalls (Messeri & Crockett, 2024), scientists can use LMs to find and synthesize literature (Van Dinter et al., 2021; Wagner et al., 2022; Fok et al., 2023; Khraisha et al., 2024) and iterate through research inquiry (Wang et al., 2023a; Morris, 2023; Liu et al., 2024). Designers can use LMs to generate and develop concept designs (Cai et al., 2023; Liu et al., 2023; Ma et al., 2023; Chong et al., 2024; Ma et al., 2024; Chen et al., 2024).

On the other hand, LMs can also aid the *refinement of ideas* – selecting from and improving upon an established existing pool of ideas. LMs can help writers by making suggested revisions (Du et al., 2022; Zhao, 2022; Mysore et al., 2023; Shu et al., 2023) and clarifying writing goals (Arnold et al., 2021; Kim et al., 2024). For scientists, LMs can facilitate revision of scientific writing (Liang et al., 2024; Radensky et al., 2024); for designers, LMs can provide feedback on (Duan et al., 2024) and annotate (Lu et al., 2024) UIs. In teaching settings, writing feedback given by LMs may be more motivating (Meyer et al., 2024) and engaging (Tanwar et al., 2024) than feedback given by other humans. Besides reviewing ideas, LMs can also curate them — for instance, by summarizing writing (Fabbri et al., 2021; Dang et al., 2022) and identifying important ideas (Lin et al., 2024).

2.2 Language Models as Critical Thinking Tools

However, one part of the thinking process is clearly missing. One does not simply go from the stimulus for ideas to figuring out how to refine them: one needs to do the actual *critical thinking*, involving reflection upon ideas, judgment, and conceptual engineering. LMs can help provide the seeds for our ideas when we don’t have any (i.e., stimulus) and help us

work through them once we've got them (i.e., refinement), but how can they help us with questioning, reorienting, analyzing, and developing ideas (i.e., critical thinking)?

There are many different definitions of critical thinking: "the propensity and skill to engage in an activity with reflective skepticism" (McPeck, 2016), "reasonable, reflective thinking that is focused on deciding what to believe or do" (Ennis, 1993), and "the development and evaluation of arguments" (Facione, 1984), among many others. Critical thinking requires many dispositions, such as seeking clear statements of questions, looking for alternatives, and being open-minded (Ennis, 1987). Critical thinking is what makes many areas of intellectual inquiry — such as history, science, and philosophy — difficult. In these areas, people must produce and work with observations that are incomplete and open to a multiplicity of framings to pursue problems with often unclear definitions of progress — a landscape which demands critical thinking. For instance, on different accounts, history requires interpreting the past with alternative (nonlinear, long-range) temporalities (Braudel, 2023), taking into account the ways in which power structures shape historical record and memory (Foucault, 1969a; Trouillot, 1995), and identifying and manipulating narrative structures (White, 1973; Gaddis, 2004). Science requires advances not only in empirical work, but also reflection upon underlying paradigms of research (Kuhn & Hawkins, 1963), epistemology (Harding, 2013), and the social and material factors that constitute scientific knowledge (Latour, 1989).

Researchers across a variety of fields have developed a rich tapestry of approaches and tools to support critical thinking and related acts. Educators develop teaching strategies to promote critical thinking (McPeck, 1990; Pithers & Soden, 2000) such as teaching and interlinking a variety of perspectives on a subject in an integrative manner (Enciso et al., 2017) and encouraging students' intellectual independence in finding answers to their questions (Langer, 1997; Raths et al., 1966). Psychologists and cognitive scientists seek to understand how cognitive mechanisms and biases inform how humans (should) develop critical thinking (Carey, 1986; Reif, 2008), emphasizing the cultivation of basic metacognitive "building blocks" of critical thinking (Pasquinelli et al., 2021) and teaching for "practical theory" (Gelder, 2005). Meanwhile, human-computer interaction (HCI) researchers explore how interactions with computer applications can facilitate critical thinking: designers can provoke experiences of discomfort (Benford et al., 2012; Halbert & Nathan, 2015); emphasize understanding over rote expression in social contexts (Kriplean et al., 2012; Sun et al., 2017); and build small "nudges" (Liao & Wang, 2022) into interfaces which "prime" (Yamamoto & Yamamoto, 2018) users towards reflective critical thinking (Bentvelzen et al., 2022); among many others. Many of these themes will be revisited in our discussion of design proposals for LMs as critical thinking tools (§5.2).

A growing body of work has explored how LMs might contribute towards critical thinking. LM-based news and media can positively affect users' willingness to think through opposing or novel viewpoints, which can be applied to combat polarization and extremism (Tanprasert et al., 2024; Zarouali et al., 2021; Shin, 2022; Wang & Tanes-Ehle, 2022; Blasiak et al., 2021). Cai et al. (2024) consider how currently "sycophantic", "servile", and "lobotomized" LMs can be used in more critical ways by challenging users' pre-existing ideas and constructively using antagonistic interactions to develop their thinking. Danry et al. (2023); Ma et al. (2023); Park & Kulkarni (2023) show how LMs can facilitate human self-reflection and improve human reasoning by asking questions instead of only answering them (as in the typical LM interaction paradigm). Xu et al. (2024) encourage critical thinking by building LM interactions using structured templates (over free-form chat). In more targeted contexts, LMs can be used to help scientific researchers critically think about their impact statements (Mukherjee et al., 2023), and to help political theorists to metacognitively reflect upon their own creative processes and judgments (Rodman, 2023).

2.3 Philosophy as Critical Thinking, Critical Thinking as Philosophy

In this paper, we focus on philosophy as a case study for critical thinking. Philosophy is concerned with critical, systematic, and reflective examination of the world. This includes understanding the basic structure of life and the world — what does it mean to exist (Aristotle, 350 BCE; Heidegger, 1927; Sartre, 1943), live (Aurelius, 180 AD), and die (Kierkegaard, 1983; Nietzsche, 1892)?; what does it mean to know something (Plato, 369BCE; Kant, 1781;

Husserl, 1931) and what are the limits of scientific knowledge (Popper, 2002; Chalmers, 2013)?; on what moral bases should we act (Aristotle, 350BCE; Spinoza, 1677), and is it even possible to determine ‘objective’ answers to moral questions (Hume, 1739; Harman & Thomson, 1996; Foucault, 1976)? Core to philosophy is “the endeavour to know how and to what extent it might be possible to think differently, instead of legitimating what is already known” (Foucault, 1976). Philosophy is for intellectual creation and engineering: Deleuze & Guattari (1991) wrote that “So long as there is a place for creating concepts, the operation that undertakes this will always be called philosophy.” In thinking about how to think, philosophy is not only about *suspicion* toward the meanings and functions of phenomena, but also *recovery* of new significances and coherence (Ricoeur, 1981).

Contrary to the image that philosophy is “done in the armchair”, isolated and impractical, philosophy has always been intertwined with other lines of inquiry. Plato engaged extensively with advanced mathematics; Aristotle contributed to early physics; Hume leaned on psychology. Philosophy has asked and continues to ask urgent, relevant questions: for instance, how are we to understand the strangeness of quantum mechanics in physics (Carnap, 1966); the relationship between consciousness (mind) and the brain (matter) (Chalmers, 2013); and “fairness” and “justice” in contexts like algorithmic discrimination (Hu, 2023), legal punishment (Alexander, 1922), and the distribution of resources (Rawls, 1971)? Indeed, researchers in every area of intellectual inquiry confront philosophical questions in their work: they might ask if a model or concept is “really real” (and how they know so), what the “nature” of their object of study is, aim to formulate normative desiderata for their theories, and so on. Therefore, we study philosophers’ views and practices in this paper both because philosophers engage extensively in critical thinking *and* because many questions which require critical thinking asked by non-philosophers often have a philosophical flavor.

3 Methods

The first author conducted interviews with 21 professional philosophers at 14 philosophy departments at doctoral universities in the United States. We contacted and selected philosophers for high diversity across area of interest (e.g., ethics, political philosophy, philosophy of science). Interviews took place online and lasted between 30 to 60 minutes, depending on interviewee availability. Interviewees were asked how they philosophize (e.g., where ideas come from, how ideas are developed, what resources are needed) and their views on LMs (e.g., can LMs ‘do’ philosophy, how might they be useful for philosophizing). These questions followed a loose script (see §B), although we asked novel follow-up questions to pursue interesting lines of inquiry raised by the interviewees’ responses. In cases where interviewees had very little or no prior exposure to LMs, they interacted live with the GPT-4 model on a philosophical topic of their choosing. We received IRB approval from our university to conduct the interviews; all interviewees confirmed their consent to participate in the study, and for their responses to inform the development of this paper. We qualitatively analyzed interview recordings and transcripts. Using an inductive approach (Thomas, 2006) and open coding (Charmaz, 2006), we identified common themes and positions (yielding §4 and §5). We refer to interviewees with a unique identifier, e.g., (P1, P2, P3) (see §A).

4 Language Models Are Not Good Critical Thinking Tools (So Far)

Many of the interviewed philosophers find LMs to be relevant and interesting, and some find them to have limited uses such as for undergraduate instruction (P1, P13, P20) or becoming acquainted with a topic (P5, P11, P12). However, none of the philosophers were convinced that current LMs can reliably and conveniently assist them in the intellectually substantive ways which require critical thinking. Philosophers described current LMs as “boring” (P2), “anodyne” (P4), “bland” (P9), and “cowardly” (P13). We discovered two broad reasons for this. First, current LMs tend to be highly neutral, detached, and non-judgmental, often commenting on ideas in abstract and decontextualized ways (§4.2). Second, current LMs tend to be servile, passive, and incurious, which is unhelpful when the user does not yet have a clear vision of what they want to accomplish, restricting the variety of intellectual interactions possible (§4.3).

4.1 How do philosophers philosophize?

A close investigation of how philosophers think through difficult philosophical questions can give us insight into the types of tools and interactions which support difficult critical thinking, and provide contrast with current LMs, which fail to perform the same function.

Where do philosophical ideas come from? Philosophers report that their ideas usually come from observing puzzles and tensions in the world, in which some aspect feels bothersome (P5, P12, P20), incomplete (P10, P14), in need of clarity (P1, P13), or outright incorrect (P3). Philosophers encounter these puzzles and tensions most commonly in open conversation with others (P1, P2, P5, P9, P19) and while reading texts — books, papers, and monographs making explicitly philosophical arguments or touching upon philosophical themes (P4, P7, P10, P12, P13, P20). These puzzles may have an intellectual or logical character: terms might not be sufficiently disambiguated, inferences may not be valid, and propositions may entail absurd conclusions (P8, P11). However, for many, these tensions are identified and drawn out by ethical motivations (P1, P8, P16, P12). Tensions might arise not primarily because a proposition is incoherent, but rather because it appears ethically problematic. For instance, the trolley problem dilemma was used to probe the differences between doing and allowing harm, with applications to bioethics, particularly abortion (Foot, 1967). Several philosophers describe being inspired by texts communicating empirical work, seeking to provide explanations for empirical observations (P1, P2, P16, P18) as well as subjecting the practices and products of the empirical sciences to critical inquiry (P2, P7, P12, P13, P18).

What do philosophers want out of their ideas? Once philosophers identify puzzles from conversations and texts, they aim to develop ideas which make progress on these puzzles. Progress is conceived of in many ways: “understand[ing] some part of the world better” (P3), working through new ways to think about problems (P17), and better understanding the current ways we think — for instance, by making implicit assumptions explicit and recognizing the implications of propositions (P7). Some philosophers describe a developed philosophical idea as a “picture” (P9, P10) which organizes subideas in a systematic way, allowing one to clearly see the main point(s). This often requires “conceptual engineering” (P6): challenging, disassembling, and rebuilding the ways in which we think.

The role of texts in philosophical development. Texts continue to actively support the philosophical development past the inception of the idea. Revisiting texts with an idea in mind can unearth new aspects of the text which comment on that idea (P9), and repeatedly consulting written ideas can be helpful for putting words to newly developed ideas (P2, P20). Because texts are static and highly accessible by many people, texts can become a shared basis for and markers in conversation with others (P9, P19). Moreover, because published texts are usually produced by people who have given a problem substantial time and thought, philosophers might approach them with more trust and charity (P4).

The role of conversation in philosophical development. Conversations with fellow philosophers are central to evaluating the coherence of ideas (P21), raising connections to other ideas and problems (P5), and collecting feedback (P3, P10). Conversations may force philosophers to explain and justify ideas they may have taken for granted (P1). Conversation helps philosophers gain confidence that their ideas are good intellectual contributions (P2, P21). Philosophers even simulate conversations in their head, taking on various positions for and against their ideas (P1, P12). Good philosophical conversation requires several conditions. The interlocutor should be charitable — genuinely listening to and working through ideas (P1, P12), and trusting (P6, P14) — but also willing to boldly push ideas forward (P3) and take intellectual risks (P18). Conversations may not be directed towards any clear goal; interlocutors must be able to “riff off each other” (P8) and be willing to operate without a preset agenda (P3, P4). This requires interlocutors to be curious about addressing problems (P21); it should be a collaborative effort, rather than a combative debate (P3, P7).

4.2 Language Models are neutral, detached, and nonjudgmental

Philosophers find intellectual value when the conversations and texts they encounter provide substantive and well-defended perspectives, but find that LMs do not do the same.

① *LMs are abstract, imprecise, and ‘skirt by’ questions.* Because philosophy is interested in clearly stating and reflecting upon ideas, philosophers often place high value on precision in language. Changes to a formulation which seem trivial to a layperson may introduce important shifts in meaning for a philosopher. Meanwhile, LMs seem as if they ‘tell the user what they want to hear’, resulting in risk-averse and hand-waving behavior which produced abstract, imprecise, and ultimately intellectually uninteresting statements (P5, P7, P15). Interviewees noted that when they brought up problems with LMs’ responses, LMs skirted around the issue, producing superficially convincing corrections without really addressing the provided issue (P1, P20). LMs are highly factually knowledgeable (P1) but fail to precisely express philosophical ideas; thus, LMs end up reinforcing the status quo rather than proposing substantive and interesting challenges (P9).

② *LM responses change too easily and don’t have ‘weight’.* Several philosophers describe how easy it is for them to talk LMs into contradictions and incoherent outputs in the same session (P4, P9). LMs make “kneejerk reactions” to user concerns and are excellent at effusively apologizing, but don’t “fully appreciate” their mistakes and the user’s comments (P14). Moreover, LM responses seem highly sensitive to trivial changes in the prompt, making some philosophers wary of using them at all (P21). The ease with which one can manipulate an LM’s output seems to reduce their trustworthiness and value as tools (P15).

③ *LM outputs don’t provide judgments.* LMs often refrain from formulating serious judgments; they try to remain neutral and ‘see all sides’, but end up presenting all sides in placid and uninteresting ways (P12, P17). They tend to refrain from discussing controversial issues (P4), which is unfortunate given that philosophy prides itself on clearly thinking about otherwise-taboo topics of controversy. As such, LMs are perceived as “cowardly”, refusing to take solid positions and, in some sense, echoing the user (P13). “It [conversations with LMs] ends up being unproductive and unsatisfying... they don’t feel like persons because their language is often so bland and impersonal, non-Socratic, generic... they’re boring” (P9).

④ *LMs don’t have memory and context.* Shared context from previous interactions with other humans serve to provide context for and situate ideas in conversation, allowing for efficiency of exploration (as already-exhausted ideas are not brought up again) (P1, P14). Because current popular LM interfaces ‘lose their memory’ of previous interactions in different sessions, LMs often produce general and decontextualized responses to user prompts (P15).

4.3 Language Models are servile, passive, and incurious

Philosophers find intellectual value when fellow philosophers develop their own lines of inquiry in conversation and texts, but find that LMs do not do the same.

① *LMs fail to be useful in open, undetermined contexts.* LMs enthusiastically make “my problem its problem” (P11), but often philosophers do not have their ‘problem’ entirely clearly thought or formulated (P5). For certain basic tasks, “you have certain success metrics in mind, so you go to [an LM]; but what about truly open-ended conversations where you don’t have success conditions already laid out?” (P7) LM answers often feel like they’ve been ‘packaged’ or return a ‘processed end result’, whereas “in the doing of philosophy, we want to be open, in service of a larger dialogue — philosophy as a process rather than as an end product” (P5). LMs don’t seem to have a drive to know the truth or care about convincing people (P2, P21) — features which interviewees note energize interactions even when there is no clearly desired product.

② *LMs restrict the variety of intellectual interaction.* The “incuriosity” of LMs severely limits possible intellectual interactions philosophers can have with it (P7). “It’s a question-answer platform. It won’t follow up with a “what do you think?” “I’m a little puzzled, how it could be?” “Oh gosh, how does it work?” You can’t have a conversation with [an LM] except one which is like an interview.” Several philosophers imagine alternative useful LM interactions in which LMs take on more intellectual risks and independent behaviors: instead of only answering questions, LMs could also ask them (P12, P17), or LMs might behave with hostility and antagonism towards users’ ideas (P6, P8, P11).

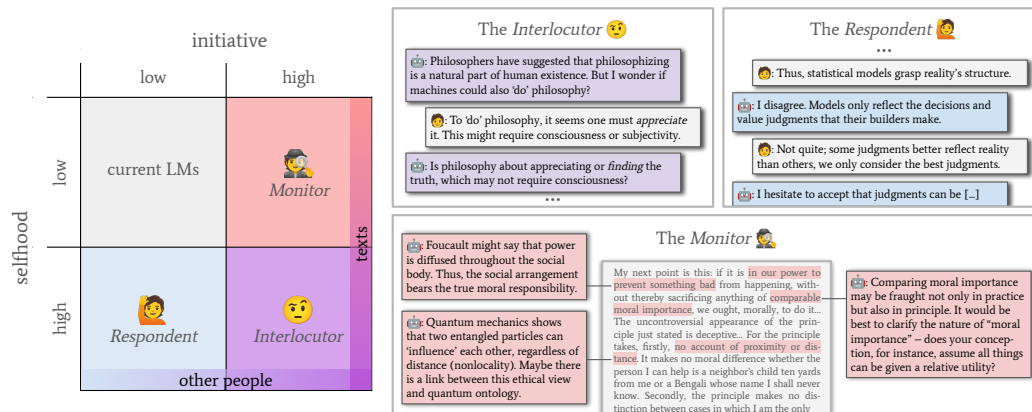


Figure 1: *Left* – The selfhood-initiative model for critical thinking tools. *Right* – Illustrative sample interactions between humans and LMs playing different roles. Other alternatives are possible. The excerpt from “The Monitor” is taken from Singer (1972).

5 Designing Language Models for Critical Thinking

Thus far, we’ve introduced the problem of critical thinking and described how current LMs fail to be good critical thinking tools for philosophers. Here, we set out a formal model to characterize and compare critical thinking tools (§5.1). This allows us to imagine new roles for LMs, inspired by what makes people and texts useful as critical thinking tools (§5.2).

5.1 The Selfhood-Initiative Model

We use the two broad reasons why LMs fail to be good critical thinking tools in §4 as the basis for the model’s two axes: current LMs have low *selfhood*, as they are neutral, detached, and nonjudgmental; they have low *initiative*, as they are servile, passive, and incurious. In particular, *selfhood* is a resource’s ability to have certain locally persistent internal states (such as perspectives, beliefs, opinions, memory) and to consistently use them as the basis for judgments. The resource’s internal states may change over time due to new knowledge and experiences, but in an intentional and logical (rather than an arbitrary and capricious) manner. Current LMs exhibit *low selfhood* (§4.2). *Initiative* is a resource’s ability to set its own intentions and goals, possibly different from its user’s, and to execute actions oriented towards those intentions. High-initiative resources are not strictly or existentially bound to their user’s directives, and may deviate from them. Current LMs exhibit *low initiative* (§4.3). These two axes form the *selfhood-initiative* model for critical thinking tools. Our model is distinct from previous models proposed for the study of critical thinking in that (a) we model *types of critical thinking tools* rather than the (*human*) *process of critical thinking* (Schön, 1987; Shneiderman, 2000, *inter alia*), and (b) we explore the *interaction* between selfhood and initiative, which have each independently been explored in some capacity by others (Cai et al., 2024; Guo et al., 2024; Hilliard et al., 2024, *inter alia*). Our model explains why philosophers find texts and other people (but not LMs) to be useful tools, and further provides a design space for LMs as critical thinking tools (§5.2).

Why do philosophers find other people and texts to be useful critical thinking tools? In the selfhood-initiative model, other people are *high-selfhood, variable-initiative* tools. People have specific backgrounds and experiences which inform their views, perspectives, and beliefs; these influence how they understand and respond. Philosophers find value in talking to other people often *because* of their selfhood; they expect that they will receive interesting judgements and comments, rather than placid neutrality. However, these people may have variable initiative, depending on the situation. In free-flowing conversation, each interlocutor may carry the conversation in some direction, whereas in a more focused conversation aimed at collecting feedback, an interlocutor may be expected to directly respond to one’s ideas and requests without their own intellectual initiative. The high selfhood of other people is

helpful because it provides particular perspectives and ways of looking into the problem space. Meanwhile, in the selfhood-initiative model, texts are *high-initiative, variable-selfhood tools*. Texts are not themselves responsive to a user's intentions (Plato, 370BCE); they express the author's attempt to fulfill their intentions, and one encounters the product of this attempt after the fact of its production. The text's exteriority from the user allows the user to reflect upon similarities and differences between their own thinking and the tool's outputs. On the other hand, the way in which texts are written can vary in the degree of selfhood they express. Informative, survey-based, and clarificatory papers tend to de-emphasize an author's perspectives and opinions, whereas more explicitly argumentative papers may center them; both can be useful to philosophers in different ways.

Why don't philosophers find current LMs to be useful critical thinking tools? In the selfhood-initiative model, current LMs are *low-selfhood, low-initiative tools*. They do not provide philosophers with particular concrete perspectives into the problem space, nor do they provide ideas sufficiently exterior to a philosopher's own thinking to allow for meaningful reflection and connections. These properties make LMs particularly useful for alternative modes of thought, such as carrying out rote and well-defined tasks and helping rewrite sentences, but not for stimulating critical thinking.

5.2 Three Roles for Language Models as Critical Thinking Tools

According to the selfhood-initiative model, good critical thinking tools should have high selfhood, high initiative, or both. From our model, we set out *three roles* of LMs for philosophy — the Interlocutor, the Monitor, and the Respondent — corresponding to the *three viable cells* in the selfhood-initiative model (high-selfhood, high-initiative; low-selfhood, high-initiative; high-selfhood, low-initiative). Implementations for these roles vary — some might be achievable with only moderate prompt engineering, whereas others might necessitate radically different user interfaces or model training methods.

The Interlocutor ◦ *high-selfhood, high-initiative*. Philosophers mention that they often get their ideas in free-flowing conversation with fellow philosophers or from reading literature that makes arguments which seem tenuous, incorrect, or incomplete (§4.1). In the terms of the selfhood-initiative model, these are *high-selfhood, high-initiative tools*. As a role for LMs, the Interlocutor would invert many of the human-AI relationships taken for granted in current LMs. Rather than attempting to remain neutral, the Interlocutor makes judgments and takes positions based on its perspectives. Rather than accommodating and affirming users' every response, the Interlocutor thinks through and challenges or disagrees with what its users say; it responds or modifies its own beliefs if users make reasonable points. Rather than remaining passive and answering user questions, the Interlocutor asks its own questions in pursuit of its 'own' interests, and refuses or redirects certain lines of inquiry in favor of others. Rather than being amnesic and detached, the Interlocutor draws upon its persistent memories and beliefs across sessions to produce ideas. The Interlocutor does not need to be strictly *antagonistic*, as explored in Cai et al. (2024); indeed, it may be charitable and polite, much like colleagues, while at the same time resisting the 'servility' and 'sycophancy' disrupted by the antagonistic paradigm.

The Monitor ◦ *low-selfhood, high-initiative*. While developing ideas, philosophers consciously or unconsciously encounter various "*decision junctures*" at which they use certain approaches or pursue certain ideas over others (P6). Many philosophers suggest that it may be important to reduce, or at least become more aware of, the choices at 'unconscious decision junctures' (P6, P2, P7). Without such awareness, philosophers may expose their ideas to imprecision ('which path did you exactly take?') and objections ('why this path and not others?'); moreover, these choices may reproduce personal and disciplinary biases, reifying metaphilosophical problems (§6.3). As a role for LMs, the Monitor acts as a 'checks and balances' on philosophizing; it is not interested in retaining self-consistency or in expressing particular points of view (low selfhood), but has high initiative to provide a variety of ideas and resources to the user. The Monitor functions similarly to survey texts which provide a 'lay of the land', illustrating different approaches and ideas to help philosophers situate their ideas, able to take all sorts of changing sides with the initiative to challenge and confront. The Monitor's suggestions may or may not be directly relevant to the philosopher's work,

but act as reference guides — to which the philosopher might think, “that’s a related idea, maybe there’s a connection here” or “that doesn’t seem directly related, but it’s good to have in mind”. Moreover, the Monitor may ask a variety of uncomfortable and unexpected methodological questions aimed at clarifying philosophers’ decisions.

The Respondent ◦ *high-selfhood, low-initiative*. As philosophers develop their ideas, they want to understand how others might react — better understanding possible misinterpretations, objections, and clarification questions which may arise (P6, P10, P12). These reactions should have high selfhood to be substantive and particular, and low initiative to remain directly focused on the user’s ideas. As a role for LMs, the Respondent adopts a specific set of beliefs and perspectives and reacts directly to the user’s ideas; it does not merely role-play or superficially caricature different positions, but should have consistent memories and beliefs which are reasonably open to change (P4) rather than dogmatically fixed. Interactions with the Respondent may inform how the philosopher formulates and presents their ideas; they may anticipate certain objections and strengthen its appeal and utility. The Respondent can also be *counterfactually* helpful: if an agent representing an unsavory position resonates with a philosopher’s argument, then that philosopher might reconsider how their argument is expressed, not only defending but also *delimiting* the scope of their argument (P6).

6 Discussion

6.1 Challenges for Language Modeling

Critical thinking can serve as another of many “north stars” in LM research, guiding what we want from LMs. Corresponding to the limitations of language models discussed in §4.2 and §4.3 are several concrete areas for further LM research. LMs will need to become more convincing agents (Andreas, 2022) which can represent specific positions and belief systems (Scherrer et al., 2023; Jin et al., 2024) ③; stay consistent with them (Chen et al., 2021; Zhao et al., 2024) ②; and commit towards and draw from long-term memory (Wang et al., 2023b) ④. In particular, LLMs will need to concretely reason about “uncommon sense” ① ②, seriously considering positions which deviate from intuitively true or correct ways of thinking about the world (Bisk et al., 2020; Ziems et al., 2023; Hendrycks et al., 2021; Pock et al., 2023). This may require rethinking how we align LMs (Ouyang et al., 2022; Sorensen et al., 2024), given that humans tend to be drawn towards confident common-sense responses (P5). LMs will need to improve their long-range planning (Hao et al., 2023) and act autonomously (Händler, 2023) ①, operating in cases where there is no clear algorithm for solving a problem (P4, P3, P8); LMs will need to take effective conceptual risks without clear immediate payoffs (P18) and reason about unsettled and open ideas (P8). To support more diverse forms of interaction beyond question answering or task execution ②, LMs will need to significantly improve in theory of mind (Jamali et al., 2023; Strachan et al., 2024). LMs need to “*understand what’s happening [in the conversation] without it being explicitly said, because.. you haven’t fully expressed it to yourself yet*” (P8), which will allow them to focus on the significant rather than irrelevant or obvious paths of inquiry in conversation (P6, P8).

6.2 Challenges for Human-AI Interaction

In addition to *modeling challenges*, there are several *interaction design challenges* when developing LMs for critical thinking. First, philosophers tend to highly value *thinking through things themselves*; many emphasize that the intellectually substantive parts of philosophy cannot be naively ‘accelerated’ (P1, P7, P14, P17). Philosophers find the process of thinking to be intrinsically valuable, even when it does not produce obvious payoffs (P3, P6, P8) — a feature common to other areas of critical thinking. Additionally, philosophers may feel that authorship of ideas requires that the ideas be ‘*mine*’, and that ‘*I*’ should be responsible for making the important intellectual judgments (P4, P10, P18). Secondly, *it can be difficult and even disruptive to put ideas into words*. Although professional philosophy is mainly formally done in language, the process of thinking through ideas can involve many other dimensions of representation and thinking (P2, P3, P4, P5). Among other challenges, philosophers cite the apparent incongruence between ideas and language as a source of significant burden

in learning how to effectively use LMs (P8, P21). This may be true for many other areas of critical thinking. Thirdly, philosophers find that *human connection is enjoyable and important*. Besides giving rise to unexpected philosophical connections and ideas (P6), conversation with another human is deeply enjoyable and fulfilling, on its own merits (P8, P21). Moreover, some philosophers feel that serious philosophical inquiry requires some kind of subjectivity or lived experience (P6, P8, P16). Therefore, LMs will need to coexist with and enrich, rather than seek to replace, the ecosystem of human and textual resources already available to philosophers and other professional critical thinkers.

6.3 LMs Help Think About and Address Metaphilosophical Problems

Throughout our interviews, we found that thinking through how LMs can serve as critical thinking tools raises many interesting metaphilosophical questions. What does it mean to ‘do’ philosophy, and who or what can ‘do’ it? How mechanical is philosophy? What is ‘thinking’? Our findings in §4.1 provide some empirical illumination for these questions. Philosophers found concretely reflecting on these questions — provoked by thinking about LMs’ role in doing philosophy — to be interesting and helpful (P1, P7, P15, P20).

However, LMs may also play an active precursory role in *addressing* metaphilosophical problems. Philosophers have articulated a host of concerns about the philosophical method and discipline: for instance, philosophers’ standards for argumentation may exclude more diverse forms of philosophical inquiry Diamond (1982); Dotson (2012), and their methods for categorizing ‘schools of thought’ (such as the analytic-continental distinction) may be counterproductive (Dolcini, 2007), reconcilable (Levy, 2003; Bell et al., 2016), and not really substantive (Mizrahi & Dickinson, 2021; Thomson, 2019). Certainly, these concerns point towards deeply entrenched sociological features of the discipline. This entrenchment is a dialectic between disciplinary structure and individual philosophers, wherein the former (materially) constrains the latter and the latter works within the lines of (and reproduces) the former. LMs might contribute towards disrupting this second direction: drawing philosophers’ attention outside the canon and across schools of thought as Interlocutors and Monitors, and representing these positions and methodologies as Respondents – possibly more approachably and accessibly than humans could. Consider Heidegger (1927)’s metaphorical carpenter: busy at work, the hammer is “ready-at-hand”, unnoticed. It is when it breaks that it becomes “present-at-hand”, noticed — an object of conscious reflection. Arguably, the philosopher must engage with ideas and methods present-, rather than ready-, at-hand (Plato, 380 BC), but the ability to engage in this way is a function of the tools and circumstances around us, and therefore often legitimately difficult (Ahmed, 2006). LMs can help, so to speak, ‘make the present-at-hand, ready-at-hand’ in a way that philosophical humans and texts cannot. Respectfully building LMs with selfhood and/or initiative into the philosopher’s material workspace – the text editor, the article viewer, and so on – can prompt ‘present-at-hand’ reflection in quiet moments and directions which a philosopher may have neglected as ready-to-hand. These small interactions, at scale, might introduce cracks into metaphilosophical edifices that philosophers would like less entrenched.

7 Conclusion: Towards Living Script

In his masterwork *Jerusalem*, Moses Mendelssohn writes that philosophy has too long prioritized a dead form of interaction, one which stifles human interaction and innovation: “We teach and instruct one another only through writings; we learn to know nature and human only from writings. We work and relax, edify and amuse ourselves through scribbling...” (Mendelssohn, 1783, 41). In response, Mendelssohn calls for a turn towards a *living script*, “arousing the mind and heart, full of meaning, continuously inspiring thought”. The living script is a way of engaging with tools that inspire and support our critical thinking; it is an ideal both for LM researchers, philosophers, and all of us — as thinkers and humans — to aspire towards. As potential technologies for reading and writing our living script, LMs can offer critical thinkers a more wide and accessible set of ways to support the development of ideas and to shape disciplinary practices and cultures. In the face of intellectual automation, it begins by saying, with John, for the rights and responsibilities to critically think: “We claim them all.”

Ethics Statement

Although exploring ‘uncommon sense’ is important for critical thinking, we acknowledge that it can also be a deeply uncomfortable and unsettling experience. Disagreement can feel awkward in many contexts in daily life, even though it may not in designated spaces: *“one of the best gifts a philosopher can give another is a good counterexample... in philosophy, we like a challenge, a pushback, for people to think that we’re wrong. That’s where philosophers thrive”* (P5). Moreover, common sense encodes certain ethical or moral norms, such as “pain is bad” and “racism is unjust”; critical thinking tools may facilitate the revisiting and challenging of these norms in apparently inappropriate ways. To be sure, there is great value in this practice. We may not only want to believe in true things but also know the right or best reasons for *why* we should believe in them (in what sense of ‘bad’ is pain *bad*? *why* is racism unjust?), since having poor reasons for a belief may undermine the belief without our knowledge. Moreover, supposedly obvious moral principles and norms can be utilized to support positions we might think to be unsavory or misguided (e.g., racism is unjust, so we should only pursue a strictly ‘colorblind’ public policy); it is difficult to identify this if one does not adopt a critical view towards the entire system. Nevertheless, LMs can serve many purposes, and being critical thinking tools is just one of them. Low-selfhood and low-initiative tools are needed to accomplish many other important tasks. Users should consent to critical interactions with LMs.

Some interviewees expressed that LMs raised difficult questions about academic integrity and authorship of ideas. It should be noted that because critical thinking tools are intended to *support* the process of thinking rather than replacing it, there is little risk of outright *plagiarism*, provided the tools are designed properly and used as intended. Nevertheless, there are interesting ethical questions about ownership of ideas with respect to involvement in their development. If a colleague’s offhand comment sparks an idea, leading to a publication, (how) should the colleague be credited? What if instead they intentionally discuss and develop an idea with you? What is an author (Foucault, 1969b)? The question of *how LMs as critical thinking tools should be credited* joins the broader existing rich discourse of how generative AI in general should be credited in intellectual production (Hullman et al., 2023; Jenkins & Lin, 2023; Simon et al., 2024; Springer, 2024, *inter alia*).

References

- Sara Ahmed. *Queer Phenomenology: Orientations, Objects, Others*. Duke University Press, Durham, NC, 2006. ISBN 978-0-8223-3914-4. URL <https://doi.org/10.1215/9780822388074>.
- Julian P. Alexander. Philosophy of punishment. *Journal of the American Institute of Criminal Law and Criminology*, 13:235, 1922.
- Jacob Andreas. Language models as agent models. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2022*, pp. 5769–5779, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.findings-emnlp.423. URL <https://aclanthology.org/2022.findings-emnlp.423>.
- Aristotle. *Metaphysics*. 350 BCE.
- Aristotle. *Nicomachean Ethics*. 350BCE.
- Kenneth C. Arnold, April M. Volzer, and Noah G. Madrid. Generative models can help writers without writing for them. In *IUI Workshops*, 2021. URL <https://ceur-ws.org/Vol-2903/IUI21WS-HAIGEN-1.pdf>.
- Marcus Aurelius. *Meditations*. 180 AD.
- Bela H. Banathy. *Designing Social Systems in a Changing World*. Springer US, 1996. ISBN 978-0-306-45251-2.
- Jeffrey A. Bell, Andrew Cutrofello, and Paul M. Livingston (eds.). *Beyond the Analytic-Continental Divide: Pluralist Philosophy in the Twenty-First Century*. Routledge, 2016. ISBN 9781138787360.
- Steve Benford, Chris Greenhalgh, Gabriella Giannachi, Brendan Walker, Joe Marshall, and Tom Rodden. Uncomfortable interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*, pp. 2005–2014, New York, NY, USA, 2012. Association for Computing Machinery. doi: 10.1145/2207676.2208347. URL <https://doi.org/10.1145/2207676.2208347>.
- Marit Bentvelzen, Paweł W. Woźniak, Pia S.F. Herbes, Evropi Stefanidi, and Jasmin Niess. Revisiting reflection in hci: Four design resources for technologies that support reflection. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 6(1), 2022. doi: 10.1145/3517233. URL <http://dx.doi.org/10.1145/3517233>.
- Bhavya Bhavya, Jinjun Xiong, and Chengxiang Zhai. Cam: A large language model-based creative analogy mining framework. In *Proceedings of the ACM Web Conference 2023*, pp. 3903–3914, New York, NY, USA, April 2023. ACM. doi: 10.1145/3543507.3587431. URL <https://doi.org/10.1145/3543507.3587431>.
- Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. PIQA: reasoning about physical commonsense in natural language, 2020. URL <https://doi.org/10.1609/aaai.v34i05.6239>.
- Kevin M. Blasiak, Marten Risius, and Sabine Matook. Conceptualising social bots for countering online extremist messages. In *ACIS 2021 Proceedings*, pp. 81, 2021. URL <https://aisel.aisnet.org/acis2021/81>.
- Fernand Braudel. *The Mediterranean and the Mediterranean World in the Age of Philip II: Volume I. The Mediterranean and the Mediterranean World in the Age of Philip II*. University of California Press, 2023. ISBN 9780520400658.
- Zana Bućinca, Maja Barbara Malaya, and Krzysztof Z. Gajos. To trust or to think: Cognitive forcing functions can reduce overreliance on ai in ai-assisted decision-making. *Proc. ACM Hum.-Comput. Interact.*, 5(CSCW1), apr 2021. doi: 10.1145/3449287. URL <https://doi.org/10.1145/3449287>.

- Alice Cai, Steven R Rick, Jennifer L Heyman, Yanxia Zhang, Alexandre Filipowicz, Matthew Hong, Matt Klenk, and Thomas Malone. Designaid: Using generative ai and semantic diversity for design inspiration. In *Proceedings of The ACM Collective Intelligence Conference, CI '23*, pp. 1–11, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400701139. doi: 10.1145/3582269.3615596. URL <https://doi.org/10.1145/3582269.3615596>.
- Alice Cai, Ian Arawjo, and Elena L. Glassman. Antagonistic ai, 2024. URL <https://arxiv.org/abs/2402.07350>.
- Alex Calderwood, Vivian Qiu, Katy Ilonka Gero, and Lydia B. Chilton. How novelists use generative language models: An exploratory user study. In Werner Geyer, Yasaman Khazaeni, and Michal Shmueli-Scheuer (eds.), *Joint Proceedings of the Workshops on Human-AI Co-Creation with Generative Models and User-Aware Conversational Agents co-located with 25th International Conference on Intelligent User Interfaces (IUI 2020)*, Cagliari, Italy, March 17, 2020, volume 2848 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2020. URL <https://ceur-ws.org/Vol-2848/HAI-GEN-Paper-3.pdf>.
- Susan Carey. Cognitive science and science education. *American psychologist*, 41(10):1123, 1986.
- Rudolf Carnap. *Philosophical Foundations of Physics*. Basic Books, Inc. Publishers, New York, first printing edition, 1966. ISBN 00008545.
- Tuhin Chakrabarty, Xurui Zhang, Smaranda Muresan, and Nanyun Peng. MERMAID: metaphor generation with symbolism and discriminative decoding. In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tür, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (eds.), *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pp. 4250–4261. Association for Computational Linguistics, 2021. doi: 10.18653/v1/2021.NAACL-MAIN.336. URL <https://doi.org/10.18653/v1/2021.naacl-main.336>.
- Tuhin Chakrabarty, Vishakh Padmakumar, Faeze Brahman, and Smaranda Muresan. Creativity support in the age of large language models: An empirical study involving emerging writers. *ArXiv*, abs/2309.12570, 2023.
- Alan F. Chalmers. *What Is This Thing Called Science?* Hackett Publishing Company, Inc., fourth edition, Sep 2013. ISBN 978-1624660382.
- Kathy Charmaz. *Constructing Grounded Theory: A Practical Guide through Qualitative Analysis*. SAGE Publications Ltd, 1 edition, 1 2006. ISBN 978-0761973539.
- Liuqing Chen, Yiyan Tsang, Qianzhi Jing, and Lingyun Sun. A llm-augmented morphological analysis approach for conceptual design. In C. Gray, E. Ciliotta Chehade, P. Hekkert, L. Forlano, P. Ciuccarelli, and P. Lloyd (eds.), *DRS2024: Boston*, Boston, USA, June 23–28 2024. Zhejiang University, Hangzhou, China; Zhejiang-Singapore Innovation and AI Joint Research Lab, Hangzhou, China. doi: 10.21606/drs.2024.605.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé de Oliveira Pinto, Jared Kaplan, Harrison Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Joshua Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code, 2021. URL <https://arxiv.org/abs/2107.03374>.

- Leah Chong, Jude Rayan, Steven Dow, Ioanna Lykourantzou, and Faez Ahmed. Cad-prompted generative models: A pathway to feasible and novel engineering designs, 2024. URL <https://arxiv.org/abs/2407.08675>.
- John Joon Young Chung, Wooseok Kim, Kang Min Yoo, Hwaran Lee, Eytan Adar, and Minsuk Chang. Talebrush: Sketching stories with generative pretrained language models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391573. doi: 10.1145/3491102.3501819. URL <https://doi.org/10.1145/3491102.3501819>.
- Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Y. Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loïc Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Searley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. No language left behind: Scaling human-centered machine translation, 2022. URL <https://doi.org/10.48550/arXiv.2207.04672>.
- Hai Dang, Karim Benharrak, Florian Lehmann, and Daniel Buschek. Beyond text generation: Supporting writers with continuous automatic text summaries. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, UIST '22, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450393201. doi: 10.1145/3526113.3545672. URL <https://doi.org/10.1145/3526113.3545672>.
- Valdemar Danry, Pat Pataranutaporn, Yaoli Mao, and Pattie Maes. Don't just tell me, ask me: Ai systems that intelligently frame explanations as questions improve human logical discernment accuracy over causal ai explanations. In *CHI '23: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 352:1–352:13, New York, NY, USA, April 2023. Association for Computing Machinery. doi: 10.1145/3544548.3580672. URL <https://doi.org/10.1145/3544548.3580672>.
- Gilles Deleuze and Félix Guattari. *What is Philosophy?* Les éditions de Minuit, France, 1996 columbia university press edition edition, 1991. ISBN 978-0231079891. English translation published in 1994 by Columbia University Press.
- Design Council. What is the framework for innovation? design council's evolved double diamond, 2019. URL <https://web.archive.org/web/20190926213512/https://www.designcouncil.org.uk/news-opinion/what-framework-innovation-design-councils-evolved-double-diamond>. Accessed: 2024-08-04.
- Cora Diamond. Anything but argument? *Philosophical Investigations*, 5(1):23–41, 1982. doi: <https://doi.org/10.1111/j.1467-9205.1982.tb00532.x>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9205.1982.tb00532.x>.
- Nevia Dolcini. The analytic/continental divide: Entities and being. *Soochow Journal of Philosophical Studies*, 16:283–302, 2007.
- Kristie Dotson. How is this paper philosophy? *Comparative Philosophy*, 3(1):03–29, 2012. ISSN 2151-6014. URL <https://www.comparativephilosophy.org>.
- Wanyu Du, Zae Myung Kim, Vipul Raheja, Dhruv Kumar, and Dongyeop Kang. Read, revise, repeat: A system demonstration for human-in-the-loop iterative text revision. In *Proceedings of the First Workshop on Intelligent and Interactive Writing Assistants (In2Writing 2022)*. Association for Computational Linguistics, 2022. doi: 10.18653/v1/2022.in2writing-1.14. URL <http://dx.doi.org/10.18653/v1/2022.in2writing-1.14>.
- Peitong Duan, Jeremy Warner, Yang Li, and Bjoern Hartmann. Generating automatic feedback on ui mockups with large language models. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA, 2024. Association

- for Computing Machinery. ISBN 9798400703300. doi: 10.1145/3613904.3642782. URL <https://doi.org/10.1145/3613904.3642782>.
- Olga Lucía Uribe Enciso, Diana Sofía Uribe Enciso, and María del Pilar Vargas Daza. Critical thinking and its importance in education: Some reflections. *Rastros Rostros*, 19(34):78–88, 2017.
- Robert H. Ennis. *A taxonomy of critical thinking dispositions and abilities*. W H Freeman/Times Books/ Henry Holt & Co, 1987.
- Robert H Ennis. Critical thinking assessment. *Theory into practice*, 32(3):179–186, 1993.
- Alexander R. Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. SummEval: Re-evaluating Summarization Evaluation. *Transactions of the Association for Computational Linguistics*, 9:391–409, 04 2021. ISSN 2307-387X. doi: 10.1162/tacl.a.00373. URL <https://doi.org/10.1162/tacl.a.00373>.
- Peter A Facione. Toward a theory of critical thinking. *Liberal Education*, 70(3):253–61, 1984.
- Giulia Di Fede, Davide Rocchesso, Steven P. Dow, and Salvatore Andolina. The idea machine: Llm-based expansion, rewriting, combination, and suggestion of ideas. pp. 623–627, 2022. doi: 10.1145/3527927.3535197. URL <https://doi.org/10.1145/3527927.3535197>.
- Raymond Fok, Hita Kambhamettu, Luca Soldaini, Jonathan Bragg, Kyle Lo, Marti A. Hearst, Andrew Head, and Daniel S. Weld. Scim: Intelligent skimming support for scientific papers. In *Proceedings of the 28th International Conference on Intelligent User Interfaces (IUI '23)*, pp. 15, New York, NY, USA, March 27–31 2023. ACM. doi: 10.1145/3581641.3584034.
- Philippa Foot. The problem of abortion and the doctrine of the double effect. *Oxford Review*, 5:5–15, 1967.
- Michel Foucault. *The Archaeology of Knowledge*. Éditions Gallimard, 1969a.
- Michel Foucault. Qu'est-ce qu'un auteur? (what is an author?). Lecture given at the Société Française de Philosophie, February 22 1969b.
- Michel Foucault. *Histoire de la sexualité: La volonté de savoir*, volume 1 of *Histoire de la sexualité*. Éditions Gallimard, Paris, 1976.
- John Lewis Gaddis. *The Landscape of History: How Historians Map the Past*. Oxford University Press, Oxford, UK, 1st edition, 2004. ISBN 978-0195171570.
- Tim van Gelder. Teaching critical thinking: Some lessons from cognitive science. *College teaching*, 53(1):41–48, 2005.
- K. Gero, Vivian Liu, and Lydia B. Chilton. Sparks: Inspiration for science writing using language models. *Proceedings of the 2022 ACM Designing Interactive Systems Conference*, 2021.
- Steven M. Goodman, Erin Buehler, Patrick Clary, Andy Coenen, Aaron Donsbach, Tiffanie N. Horne, Michal Lahav, Robert MacDonald, Rain Breaw Michaels, Ajit Narayanan, Mahima Pushkarna, Joel Riley, Alex Santana, Lei Shi, Rachel Sweeney, Phil Weaver, Ann Yuan, and Meredith Ringel Morris. Lampost: Design and evaluation of an ai-assisted email writing prototype for adults with dyslexia. In *The 24th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '22)*, pp. 17, Athens, Greece, 2022. ACM. doi: 10.1145/3517428.3544819.
- Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. Large language model based multi-agents: A survey of progress and challenges, 2024. URL <https://arxiv.org/abs/2402.01680>.

- Helen Halbert and Lisa P. Nathan. Designing for discomfort: Supporting critical reflection through interactive tools. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing, CSCW '15*, pp. 349–360, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450329224. doi: 10.1145/2675133.2675162. URL <https://doi.org/10.1145/2675133.2675162>.
- Thorsten Händler. Balancing autonomy and alignment: A multi-dimensional taxonomy for autonomous llm-powered multi-agent architectures, 2023. URL <https://doi.org/10.48550/arXiv.2310.03659>.
- Shibo Hao, Yi Gu, Haodi Ma, Joshua Hong, Zhen Wang, Daisy Wang, and Zhiting Hu. Reasoning with language model is planning with world model, December 2023. URL <https://aclanthology.org/2023.emnlp-main.507>.
- Sandra Harding. Rethinking standpoint epistemology: What is “strong objectivity”? In *Feminist epistemologies*, pp. 49–82. Routledge, 2013.
- Gilbert Harman and Judith Thomson. *Moral Relativism and Moral Objectivity*. Wiley-Blackwell, 1st edition, Jan 1996. ISBN 978-0631192114.
- Martin Heidegger. *Being and Time*. 1927. Original title: Sein und Zeit. Translated by John Macquarrie and Edward Robinson (1962), Joan Stambaugh (1996).
- Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob Steinhardt. Aligning AI with shared human values, 2021. URL https://openreview.net/forum?id=dNy_RKzJacY.
- Airlie Hilliard, Cristian Muñoz, Zekun Wu, and Adriano Soares Koshiyama. Eliciting personality traits in large language models. *CoRR*, abs/2402.08341, 2024. doi: 10.48550/ARXIV.2402.08341. URL <https://doi.org/10.48550/arXiv.2402.08341>.
- Lily Hu. What is race? in algorithmic discrimination on the basis of race? *Journal of Moral Philosophy*, 21(1-2):1–26, 2023. doi: 10.1163/17455243-20234369.
- Jessica Hullman, Ari Holtzman, and Andrew Gelman. Artificial intelligence and aesthetic judgment, 2023. URL <https://arxiv.org/abs/2309.12338>.
- David Hume. *A Treatise of Human Nature*. John Noon, London, 1739.
- Edmund Husserl. *Cartesian Meditations: An Introduction to Phenomenology*. Armand Collin, 1931.
- Aldous Huxley. *Brave New World*. Harper & Brothers, New York, 1932. ISBN 978-0-06-085052-4.
- Mohsen Jamali, Ziv M. Williams, and Jing Cai. Unveiling theory of mind in large language models: A parallel to single neurons in the human brain, 2023. URL <https://doi.org/10.48550/arXiv.2309.01660>.
- Ryan Jenkins and Patrick Lin. Ai-assisted authorship: How to assign credit in synthetic scholarship. Report, Ethics + Emerging Sciences Group, 2023.
- Mingyu Jin, Beichen Wang, Zhaoqian Xue, Suiyuan Zhu, Wenyue Hua, Hua Tang, Kai Mei, Mengnan Du, and Yongfeng Zhang. What if llms have different world views: Simulating alien civilizations with llm-based agents, 2024. URL <https://doi.org/10.48550/arXiv.2402.13184>.
- Immanuel Kant. *Critique of Pure Reason*. 1781.
- Qusai Khraisha, Sophie Put, Johanna Kappenberg, Azza Warraitch, and Kristin Hadfield. Can large language models replace humans in systematic reviews? evaluating gpt-4’s efficacy in screening and extracting data from peer-reviewed and grey literature in multiple languages. *Research Synthesis Methods*, 15(4):616–626, July 2024. doi: 10.1002/jrsm.1715. Epub 2024 Mar 14.

- Søren Kierkegaard. *The Sickness Unto Death: A Christian Psychological Exposition For Upbuilding And Awakening*, volume 19 of *Kierkegaard's Writings*. Princeton University Press, Nov 1983. ISBN 9780691020280.
- Jeongyeon Kim, Sangho Suh, Lydia B. Chilton, and Haijun Xia. Metaphorian: Leveraging large language models to support extended metaphor creation for science writing. In *Designing Interactive Systems Conference (DIS '23)*, pp. 21, New York, NY, USA, 2023. ACM. doi: 10.1145/3563657.3595996. URL <https://doi.org/10.1145/3563657.3595996>.
- Jiho Kim, Ray C. Flanagan, Noelle E. Haviland, Zeai Sun, Souad N. Yakubu, Edom A. Maru, and Kenneth C. Arnold. Towards full authorship with AI: supporting revision with ai-generated views, 2024. URL <https://ceur-ws.org/Vol-3660/paper17.pdf>.
- Travis Kriplean, Michael Toomim, Jonathan Morgan, Alan Borning, and Amy J. Ko. Is this what you meant? promoting listening on the web with reflect. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '12*, pp. 1559–1568, New York, NY, USA, 2012. Association for Computing Machinery. ISBN 9781450310154. doi: 10.1145/2207676.2208621. URL <https://doi-org.offcampus.lib.washington.edu/10.1145/2207676.2208621>.
- Thomas S. Kuhn and David Hawkins. The structure of scientific revolutions. *American Journal of Physics*, 31:554–555, 1963.
- Michelle S. Lam, Janice Teoh, James A. Landay, Jeffrey Heer, and Michael S. Bernstein. Concept induction: Analyzing unstructured text with high-level concepts using lloom. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM, 2024. ISBN 979-8-4007-0330-0/24/05. doi: 10.1145/3613904.3642830. URL <https://doi.org/10.1145/3613904.3642830>.
- Ellen J. Langer. *The Power of Mindful Learning*. Addison-Wesley, New York, 1997.
- Bruno Latour. Science in action : how to follow scientists and engineers through society. *Contemporary Sociology*, 18:788, 1989.
- Neil Levy. Analytic and continental philosophy: Explaining the differences. *Metaphilosophy*, 34(3):284–304, 2003. ISSN 00261068, 14679973. URL <http://www.jstor.org/stable/24439383>.
- Weixin Liang, Yuhui Zhang, Hancheng Cao, Binglu Wang, Daisy Yi Ding, Xinyu Yang, Kailas Vodrahalli, Siyu He, Daniel Scott Smith, Yian Yin, Daniel A. McFarland, and James Zou. Can large language models provide useful feedback on research papers? a large-scale empirical analysis. *NEJM AI*, 1(8):AIoa2400196, 2024. doi: 10.1056/AIoa2400196. URL <https://ai.nejm.org/doi/abs/10.1056/AIoa2400196>.
- Jingxian Liao and Hao-Chuan Wang. Nudge for reflective mind: Understanding how accessing peer concept mapping and commenting affects reflection of high-stakes information. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems, CHI EA '22*, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391566. doi: 10.1145/3491101.3519815. URL <https://doi.org/10.1145/3491101.3519815>.
- Susan Lin, Jeremy Warner, J.D. Zamfirescu-Pereira, Matthew G Lee, Sauhard Jain, Shanqing Cai, Piyawat Lertvittayakumjorn, Michael Xuelin Huang, Shumin Zhai, Bjoern Hartmann, and Can Liu. Rambler: Supporting writing with speech via llm-assisted gist manipulation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems, CHI '24*, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400703300. doi: 10.1145/3613904.3642217. URL <https://doi.org/10.1145/3613904.3642217>.
- Vivian Liu, Jo Vermeulen, George Fitzmaurice, and Justin Matejka. 3dall-e: Integrating text-to-image ai in 3d design workflows. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference, DIS '23*, pp. 1955–1977, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450398930. doi: 10.1145/3563657.3596098. URL <https://doi.org/10.1145/3563657.3596098>.

- Yiren Liu, Si Chen, Haocong Cheng, Mengxia Yu, Xiao Ran, Andrew Mo, Yiliu Tang, and Yun Huang. How ai processing delays foster creativity: Exploring research question co-creation with an llm-based agent. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, pp. 25, New York, NY, USA, 2024. ACM. doi: 10.1145/3613904.3642698. URL <https://doi.org/10.1145/3613904.3642698>.
- Yuwen Lu, Ziang Tong, Qinyi Zhao, Yewon Oh, Bryan Wang, and Toby Jia-Jun Li. Flowy: Supporting ux design decisions through ai-driven pattern annotation in multi-screen user flows, 2024. URL <https://arxiv.org/abs/2406.16177>.
- Jenny Ma, Karthik Sreedhar, Vivian Liu, Sitong Wang, Pedro Alejandro Perez, and Lydia B. Chilton. Didup: Dynamic iterative development for ui prototyping, 2024. URL <https://arxiv.org/abs/2407.08474>.
- Kevin Ma, Daniele Grandi, Christopher McComb, and Kosa Goucher-Lambert. Conceptual design generation using large language models, 2023. URL <https://doi.org/10.48550/arXiv.2306.01779>.
- John E. McPeck. *Teaching Critical Thinking: Dialogue and Dialectic*. Routledge, 1990. ISBN 0415902258, 9780415902250.
- John E. McPeck. *Critical thinking and education*. Routledge, 2016.
- Moses Mendelssohn. *Jerusalem, or on Religious Power and Judaism*. Friedrich Maurer, Berlin, 1783.
- Lisa Messeri and M. J. Crockett. Artificial intelligence and illusions of understanding in scientific research. *Nature*, 627(8002):49–58, March 1 2024. doi: 10.1038/s41586-024-07146-0. URL <https://doi.org/10.1038/s41586-024-07146-0>.
- Jennifer Meyer, Thorben Jansen, Ronja Schiller, Lucas W. Liebenow, Marlene Steinbach, Andrea Horbach, and Johanna Fleckenstein. Using llms to bring evidence-based feedback into the classroom: Ai-generated feedback increases secondary students' text revision, motivation, and positive emotions. *Computers and Education: Artificial Intelligence*, 6: 100199, 2024. ISSN 2666-920X. doi: <https://doi.org/10.1016/j.caeai.2023.100199>. URL <https://www.sciencedirect.com/science/article/pii/S2666920X23000784>.
- Piotr Mirowski, Kory W. Mathewson, Jaylen Pittman, and Richard Evans. Co-writing screenplays and theatre scripts with language models: Evaluation by industry professionals. pp. 355:1–355:34, 2023. doi: 10.1145/3544548.3581225. URL <https://doi.org/10.1145/3544548.3581225>.
- Moti Mizrahi and Mike Dickinson. The analytic-continental divide in philosophical practice: An empirical study. *Metaphilosophy*, 52(5):668–680, 2021. doi: 10.1111/meta.12519.
- Meredith Ringel Morris. Scientists' perspectives on the potential for generative ai in their fields, 2023. URL <https://arxiv.org/abs/2304.01420>.
- Anwasha Mukherjee, Vagner Figuerêdo de Santana, and Alexis T. Baria. Impactbot: Chatbot leveraging language models to automate feedback and promote critical thinking around impact statements. In Albrecht Schmidt, Kaisa Väänänen, Tesh Goyal, Per Ola Kristensson, and Anicia Peters (eds.), *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems, CHI EA 2023, Hamburg, Germany, April 23-28, 2023*, pp. 388:1–388:8. ACM, 2023. doi: 10.1145/3544549.3573844. URL <https://doi.org/10.1145/3544549.3573844>.
- Sheshera Mysore, Zhuoran Lu, Mengting Wan, Longqi Yang, Steve Menezes, Tina Baghaee, Emmanuel Barajas Gonzalez, Jennifer Neville, and Tara Safavi. PEARL: personalizing large language model writing assistants with generation-calibrated retrievers. *CoRR*, abs/2311.09180, 2023. doi: 10.48550/ARXIV.2311.09180. URL <https://doi.org/10.48550/arXiv.2311.09180>.
- Friedrich Nietzsche. *Also sprach Zarathustra: Ein Buch für Alle und Keinen*. Ernst Schmeitzner, Germany, 1892. Published in parts between 1883 and 1892.

- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/b1efde53be364a73914f58805a001731-Abstract-Conference.html.
- Soya Park and Chinmay Kulkarni. Thinking assistants: Llm-based conversational assistants that help users think by asking rather than answering. *CoRR*, abs/2312.06024, 2023. doi: 10.48550/ARXIV.2312.06024. URL <https://doi.org/10.48550/arXiv.2312.06024>.
- Elena Pasquinelli, Mathieu Farina, Audrey Bedel, and Roberto Casati. Naturalizing critical thinking: consequences for education, blueprint for future research in cognitive science. *Mind, Brain, and Education*, 15(2):168–176, 2021.
- Robert T Pithers and Rebecca Soden. Critical thinking in education: A review. *Educational research*, 42(3):237–249, 2000.
- Plato. *Theaetetus*. 369BCE.
- Plato. *Phaedrus*. 370BCE.
- Plato. *The Republic*. 380 BC.
- Mark Pock, Andre Ye, and Jared Moore. Llms grasp morality in concept, 2023. URL <https://arxiv.org/abs/2311.02294>.
- Karl Popper. *The Logic of Scientific Discovery*. Routledge Classics, 2nd edition, Feb 2002. ISBN 978-0415278447.
- Marissa Radensky, Daniel S. Weld, Joseph Chee Chang, Pao Siangliulue, and Jonathan Bragg. Let’s get to the point: Llm-supported planning, drafting, and revising of research-paper blog posts, 2024. URL <https://arxiv.org/abs/2406.10370>.
- LE Raths, S Wasserman, A Jonas, and A Rothstein. Teaching for critical thinking: Theory and application. *Columbus, OH: Charles-Merrill*, 1966.
- John Rawls. *A Theory of Justice: Original Edition*. Harvard University Press, 1971. ISBN 9780674880108. URL <http://www.jstor.org/stable/j.ctvjf9z6v>.
- Jude Rayan, Dhruv Kanetkar, Nicole Gong, Yüewen Yang, Srishti Palani, Haijun Xia, and Steven P. Dow. Exploring the potential for generative ai-based conversational cues for real-time collaborative ideation. In *Creativity and Cognition (C&C ’24)*, pp. 15, New York, NY, USA, June 23–26 2024. ACM. doi: 10.1145/3635636.3656184.
- Frederick Reif. *Applying cognitive science to education: Thinking and learning in scientific and other complex domains*. MIT press, 2008.
- Paul Ricoeur. *Hermeneutics and the Human Sciences: Essays on Language, Action and Interpretation*. Cambridge University Press, Cambridge, 1981. ISBN 0521280028.
- Emma Rodman. On political theory and large language models. *Political Theory*, 2023. doi: 10.1177/0090591723120082. URL <https://doi.org/10.1177/0090591723120082>.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton-Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. Code llama: Open foundation models for code. *CoRR*, abs/2308.12950, 2023. doi: 10.48550/ARXIV.2308.12950. URL <https://doi.org/10.48550/arXiv.2308.12950>.

- Jean-Paul Sartre. *Being and Nothingness: An Essay on Phenomenological Ontology*. Éditions Gallimard, 1943.
- Nino Scherrer, Claudia Shi, Amir Feder, and David M. Blei. Evaluating the moral beliefs encoded in llms. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/a2cf225ba392627529efef14dc857e22-Abstract-Conference.html.
- Oliver Schmitt and Daniel Buschek. Characterchat: Supporting the creation of fictional characters through conversation and progressive manifestation with a chatbot. In Corina Sas, Neil A. M. Maiden, Brian P. Bailey, Celine Latulipe, and Ellen Yi-Luen Do (eds.), *Proceedings of the 13th ACM SIGCHI Conference on Creativity and Cognition, C&C 2021, Virtual Event / Venice, Italy, June 22-23, 2021*, pp. 7:1–7:10. ACM, 2021. doi: 10.1145/3450741.3465253. URL <https://doi.org/10.1145/3450741.3465253>.
- Eric Schwitzgebel, David Schwitzgebel, and Anna Strasser. Creating a large language model of a philosopher. *CoRR*, abs/2302.01339, 2023. doi: 10.48550/ARXIV.2302.01339. URL <https://doi.org/10.48550/arXiv.2302.01339>.
- Donald A. Schön. *Educating the Reflective Practitioner: Toward a New Design for Teaching and Learning in the Professions*. Jossey-Bass, San Francisco, CA, 1987. ISBN 978-1555420253.
- Orit Shaer, Angelora Cooper, Osnat Mokryn, Andrew L Kun, and Hagit Ben Shoshan. Ai-augmented brainwriting: Investigating the use of llms in group ideation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems, CHI '24, New York, NY, USA, 2024*. Association for Computing Machinery. ISBN 9798400703300. doi: 10.1145/3613904.3642414. URL <https://doi.org/10.1145/3613904.3642414>.
- Donghee Shin. The perception of humanness in conversational journalism: An algorithmic information-processing perspective. *New Media & Society*, 24(12):2680–2704, 2022. doi: 10.1177/1461444821993801. URL <https://doi.org/10.1177/1461444821993801>.
- Ben Shneiderman. Creating creativity: user interfaces for supporting innovation. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 7(1):114–138, 2000.
- Lei Shu, Liangchen Luo, Jayakumar Hoskere, Yun Zhu, Yinxiao Liu, Simon Tong, Jindong Chen, and Lei Meng. Rewritelms: An instruction-tuned large language model for text rewriting, 2023.
- J. Simon, G. Rieder, and J. Branford. The philosophy and ethics of ai: Conceptual, empirical, and technological investigations into values. *DISO*, 3(10), 2024. doi: 10.1007/s44206-024-00094-2. URL <https://doi.org/10.1007/s44206-024-00094-2>.
- Peter Singer. Famine, affluence, and morality. *Philosophy and Public Affairs*, 1(3):229–243, 1972. ISSN 00483915, 10884963. URL <http://www.jstor.org/stable/2265052>.
- Taylor Sorensen, Jared Moore, Jillian Fisher, Mitchell L. Gordon, Niloofar Miresghallah, Christopher Michael Rytting, Andre Ye, Liwei Jiang, Ximing Lu, Nouha Dziri, Tim Althoff, and Yejin Choi. A roadmap to pluralistic alignment, 2024. URL <https://doi.org/10.48550/arXiv.2402.05070>.
- Baruch Spinoza. *Ethics, Demonstrated in Geometrical Order*. Posthumous, 1677.
- Springer. Artificial intelligence (ai) - editorial policies. <https://www.springer.com/gp/editorial-policies/artificial-intelligence--ai-/25428500>, 2024. Accessed: 2024-07-29.
- James W. A. Strachan, Dalila Albergo, Giulia Borghini, Oriana Pansardi, Eugenio Scaliti, Saurabh Gupta, Krati Saxena, Alessandro Rufo, Stefano Panzeri, Guido Manzi, Michael S. A. Graziano, and Cristina Becchio. Testing theory of mind in large language models and humans. *Nature Human Behaviour*, 8(7):1285–1295, 2024. doi: 10.1038/s41562-024-01882-z. URL <https://doi.org/10.1038/s41562-024-01882-z>.

- Na Sun, Chien Wen (Tina) Yuan, Mary Beth Rosson, Yu Wu, and Jack M. Carroll. Critical thinking in collaboration: Talk less, perceive more. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, CHI EA '17, pp. 2944–2950, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450346566. doi: 10.1145/3027063.3053250. URL <https://doi-org.offcampus.lib.washington.edu/10.1145/3027063.3053250>.
- Thitaree Tanprasert, Sidney S Fels, Luanne Sinnamon, and Dongwook Yoon. Debate chatbots to facilitate critical thinking on youtube: Social identity and conversational style make a difference. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400703300. doi: 10.1145/3613904.3642513. URL <https://doi.org/10.1145/3613904.3642513>.
- Henansh Tanwar, Kunal Shrivastva, Rahul Singh, and Dhruv Kumar. Opinebot: Class feedback reimaged using a conversational LLM. *CoRR*, abs/2401.15589, 2024. doi: 10.48550/ARXIV.2401.15589. URL <https://doi.org/10.48550/arXiv.2401.15589>.
- David R. Thomas. A general inductive approach for analyzing qualitative evaluation data. *American Journal of Evaluation*, 27(2):237–246, 2006. doi: 10.1177/1098214005283748.
- Iain D. Thomson. *Rethinking the Analytic/Continental Divide*, pp. 569–589. Cambridge University Press, 2019.
- Michel-Rolph Trouillot. *Silencing the Past: Power and the Production of History*. Beacon Press books. Beacon Press, 1995. ISBN 9780807043110.
- Raymon Van Dinter, Bedir Tekinerdogan, and Cagatay Catal. Automation of systematic literature reviews: A systematic literature review. *Information and Software Technology*, 136: 106589, 2021.
- Gerit Wagner, Roman Lukyanenko, and Guy Paré. Artificial intelligence and the conduct of literature reviews. *Journal of Information Technology*, 37(2):209–226, 2022.
- Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal Chandak, Shengchao Liu, Peter Van Katwyk, Andreea Deac, Anima Anandkumar, Karianne Bergen, Carla P. Gomes, Shirley Ho, Pushmeet Kohli, Joan Lasenby, Jure Leskovec, Tie-Yan Liu, Arjun Manrai, Debora Marks, Bharath Ramsundar, Le Song, Jimeng Sun, Jian Tang, Petar Veličković, Max Welling, Linfeng Zhang, Connor W. Coley, Yoshua Bengio, and Marinka Zitnik. Scientific discovery in the age of artificial intelligence. *Nature*, 620(7972):47–60, August 1 2023a. doi: 10.1038/s41586-023-06221-2. URL <https://doi.org/10.1038/s41586-023-06221-2>.
- Jinping Wang and Zeynep Tanes-Ehle. Examining the effects of conversational chatbots on changing conspiracy beliefs about science: The paradox of interactivity. *Journal of Broadcasting & Electronic Media*, 67(1):68–89, 2022. doi: 10.1080/08838151.2022.2153842. URL <https://doi.org/10.1080/08838151.2022.2153842>.
- Weizhi Wang, Li Dong, Hao Cheng, Xiaodong Liu, Xifeng Yan, Jianfeng Gao, and Furu Wei. Augmenting language models with long-term memory. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023b. URL http://papers.nips.cc/paper_files/paper/2023/hash/ebd82705f44793b6f9ade5a669d0f0bf-Abstract-Conference.html.
- Hayden White. *Metahistory: The Historical Imagination in Nineteenth-century Europe*. Johns Hopkins University, 1973.
- Xiaotong (Tone) Xu, Jiayu Yin, Catherine Gu, Jenny Mar, Sydney Zhang, Jane L. E, and Steven P. Dow. Jamplate: Exploring llm-enhanced templates for idea reflection. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*, IUI '24, pp. 907–921, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400705083. doi: 10.1145/3640543.3645196. URL <https://doi.org/10.1145/3640543.3645196>.

- Yusuke Yamamoto and Takehiro Yamamoto. Query priming for promoting critical thinking in web search. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval, CHIIR '18*, pp. 12–21, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450349253. doi: 10.1145/3176349.3176377. URL <https://doi.org/10.1145/3176349.3176377>.
- Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. Wordcraft: Story writing with large language models. In *Proceedings of the 27th International Conference on Intelligent User Interfaces, IUI '22*, pp. 841–852, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391443. doi: 10.1145/3490099.3511105. URL <https://doi.org/10.1145/3490099.3511105>.
- Brahim Zarouali, Mykola Makhortykh, Mariella Bastian, and Theo Araujo. Overcoming polarization with chatbot news? investigating the impact of news content containing opposing views on agreement and credibility. *European Journal of Communication*, 36(1):53–68, 2021. doi: 10.1177/0267323120940908. URL <https://doi.org/10.1177/0267323120940908>.
- Xin Zhao. Leveraging artificial intelligence (ai) technology for english writing: Introducing wordtune as a digital writing assistant for efl writers. *RELC Journal*, 54:890 – 894, 2022.
- Yukun Zhao, Lingyong Yan, Weiwei Sun, Guoliang Xing, Shuaiqiang Wang, Chong Meng, Zhicong Cheng, Zhaochun Ren, and Dawei Yin. Improving the robustness of large language models via consistency alignment. In Nicoletta Calzolari, Min-Yen Kan, Véronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue (eds.), *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC/COLING 2024, 20-25 May, 2024, Torino, Italy*, pp. 8931–8941. ELRA and ICCL, 2024. URL <https://aclanthology.org/2024.lrec-main.782>.
- Caleb Ziemis, Jane Dwivedi-Yu, Yi-Chia Wang, Alon Y. Halevy, and Diyi Yang. Normbank: A knowledge bank of situational social norms. In Anna Rogers, Jordan L. Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pp. 7756–7776. Association for Computational Linguistics, 2023. doi: 10.18653/V1/2023.ACL-LONG.429. URL <https://doi.org/10.18653/v1/2023.acl-long.429>.

A Interviewee Information Sheet

Table 1 provides high-level information about each interviewee which may be relevant to interpreting and contextualizing their views. The *General Interest(s)* feature describes the broad fields that the interviewees work in. The *Notable Specific Interest(s)* feature describes any specific topics in the field(s) mentioned in the *General Interest(s)* feature that the interviewees focus their work on. This feature is not exclusive, meaning that interviewees may also work on other topics outside of the specific interests. If the value for this feature is blank, then the interviewee's work is sufficiently characterized by the value in the *General Interest(s)* feature. The *Experience with LMs* feature describes three levels of experience with using LMs: little to none, limited, and extensive. If interviewees have either limited or extensive experience with using LMs, the *Uses of LMs* feature describes their primary use: for teaching (e.g., using LMs to teach material, trying to understand features of LM-generated student submissions), for personal use (e.g., to improve productivity, for entertainment), for exploration (i.e., playing around with the LM out of curiosity to understand the technology better), and for research (i.e., their research is on LMs). Note that the following interviewees have published at least one article on some aspect of AI: (P5, P6, P13, P14).

ID	Title	General Interest(s)	Notable Specific Interest(s)	Experience with LMs	Use of LMs
P1	Associate Professor	Ethics, Political Philosophy	Bioethics, Feminist Ethics	Limited	For teaching
P2	Associate Professor	Philosophy of Science	Philosophy of Biology	Limited	For exploration
P3	Professor	Ethics, Aesthetics	Meta-ethics	Limited	For teaching
P4	Professor	Ethics, Political Philosophy		Limited	For personal use
P5	Assistant Professor	Ethics	Virtue ethics	Limited	For teaching
P6	Assistant Professor	Ethics, Political Philosophy	Philosophy of Technology, AI	Extensive	For research
P7	Assistant Professor	Philosophy of Science	Philosophy of Physics	Extensive	For personal use
P8	Associate Professor	History of Philosophy	German philosophy	Limited	For personal use
P9	Professor	Philosophy of Science		Extensive	For exploration
P10	Associate Professor	Ethics, History of Philosophy	Philosophy of Technology	Little to None	
P11	Professor	Philosophy of Science	Philosophy of Statistics	Little to None	
P12	Professor	Philosophy of Science	Psychology	Limited	For exploration
P13	Associate Professor	Philosophy of Science	Philosophy of Biology	Limited	For class
P14	Professor	Logic, Philosophy of Mind	Semantics, Linguistics	Extensive	For exploration
P15	Assistant Professor	Aesthetics	Value theory, Literature	Limited	For exploration
P16	Professor	Ethics, Political Philosophy	Public and Global Policy	Limited	For exploration
P17	Teaching Professor	Pedagogy, Epistemology		Extensive	For exploration
P18	Associate Professor	Philosophy of Science	Philosophy of Physics	Limited	For class
P19	Assistant Professor	Ethics	Moral psychology	Limited	For personal use
P20	Professor	History of Philosophy		Little to None	
P21	Associate Professor	Ethics	Bioethics	Little to None	

Table 1: Interviewee information.

B Interview Questions and Guidelines

1. Meta-philosophy
 - (a) What is philosophy? Why do you go about doing philosophy? What aims do you have?
 - (b) What drives the 'doing' of philosophy? What is the role of personal motivations, subjective experience, and aesthetic judgements?
 - (c) Who or what can 'do' philosophy? For instance, can LLMs 'do' philosophy?
 - (d) What makes doing philosophy 'difficult' / nontrivial?
 - (e) How does philosophy distinguish its products from those of other disciplines?
2. The philosophical process
 - (a) How do you go from no idea to a spark of an idea / an unrefined idea?
 - (b) How do you develop and refine philosophical ideas? What moves have to happen?
 - (c) How mechanical / creative is the process of doing philosophy?
 - (d) What is the relationship between texts / textual methods and philosophy? Does philosophizing, to some extent, operate 'above' language in ideas / thoughts?
 - (e) What is the role of conversation in the doing of philosophy? What are some of its challenges?
 - (f) What makes for a good interlocutor, and what makes for a good conversation?
3. Language Models for philosophy
 - (a) What roles can language models play in the development of philosophy?
 - (b) What do language models need to be better in the development of philosophy?
 - (c) What are some of the opportunities and strengths for language models in philosophy?
 - (d) What are some of the risks and weaknesses for language models in philosophy?
 - (e) Would you use language models in intellectually substantive ways currently? What about in the future, with plausible improvements?