

Unsocial Intelligence: An Investigation of the Assumptions of AGI Discourse

Borhane Blili-Hamelin^{*1}, Leif Hancox-Li^{*2}, Andrew Smart³

¹AI Risk and Vulnerability Alliance

borhane@avidml.org

²vijil

leif@vijil.ai

³Google Research

andrewsmart@google.com

Abstract

Dreams of machines rivaling human intelligence have shaped the field of AI since its inception. Yet, the very meaning of human-level AI or artificial general intelligence (AGI) remains elusive and contested. Definitions of AGI embrace a diverse range of incompatible values and assumptions. Contending with the fractured worldviews of AGI discourse is vital for critiques that pursue different values and futures. To that end, we provide a taxonomy of AGI definitions, laying the ground for examining the key social, political, and ethical assumptions they make. We highlight instances in which these definitions frame AGI or human-level AI as a technical topic and expose the value-laden choices being implicitly made. Drawing on feminist, STS, and social science scholarship on the political and social character of intelligence in both humans and machines, we propose contextual, democratic, and participatory paths to imagining future forms of machine intelligence. The development of future forms of AI must involve explicit attention to the values it encodes, the people it includes or excludes, and a commitment to epistemic justice.

1 Introduction

There is no agreed-upon definition of artificial general intelligence (AGI) (Morris et al. 2023; Mitchell 2024). Yet from influential AI companies (DeepMind 2022; OpenAI 2018; Meaker 2023; Ingram 2024) and AI researchers (Agiuera y Arcas and Norvig 2023; OpenAI 2018; Morris et al. 2023; Chollet 2019; Bubeck et al. 2023) to the increasingly public worry (Future of Life Institute 2023; Center for AI Safety 2024) about existential risks (Bostrom 2014; McLean et al. 2023; Naudé and Barten 2023; Noy and Uher 2022), AGI and human-level AI have become one of the dominant ways to imagine the future potential of AI. At the same time, AGI has had its own skeptics (LeCun 2022; Toews 2022; Heaven 2020; Gebru and Torres 2024) who risk taking the target of their critiques to be more homogenous and congruent than it actually is.

The lack of homogeneity in current conceptions of AGI is not a bug. It is a feature of the underlying topic: what might it mean for machines to have intelligence that rivals

human intelligence? Substantive disagreements come with the value-laden character of both intelligence and of AI as a technology. By value-laden, we mean that political, social, and ethical values can and should shape conceptions of intelligence, technology, and their intersection. (Blili-Hamelin and Hancox-Li 2023) We lack consensus on the definition of intelligence because of conflicting values. The same applies to the current lack of consensus on AGI.

Ruha Benjamin warns that “a narrow definition of what even counts as technology or intelligence” threatens the ability of communities to imagine worlds worth building. (Benjamin 2024) What forms of intelligence and technology are worth imagining are political and social questions. However, current approaches to AGI risk mistaking these questions for technical questions. We lay the groundwork for examining these potential mistakes by analyzing and classifying contingent assumptions underlying different definitions of AGI. Contending with the fractured worldviews among conceptions of AGI plants the seeds for resisting harmfully narrow conceptions of machine intelligence.

Our paper proceeds as follows. We begin by providing a framework for thinking of AGI as inheriting the value-laden features of both intelligence and technology (Section 2). We then investigate the value-laden, contingent choices made by influential accounts of AGI and human-level AI (Section 3). Rather than framing these choices as inherently misguided, we see them as opportunities to investigate the limited but heterogeneous range of questions currently being asked about would-be human-level AI. Finally, we sketch pathways for more contextual, democratically legitimate, and participatory perspectives on what forms of machine intelligence are worth imagining (Sections 4 and 5).

Our investigation is not limited to accounts using the exact phrase “artificial general intelligence”. Like Morris et al. 2023, our topic is the long history of treating “[a]chieving human-level “intelligence” as the “north-star goal” of the AI field (Morris et al. 2023), dating at least as far back as the 1955 Dartmouth AI Conference (McCarthy et al. 1955). Discussions of “human-level AI”, “general AI”, or “AI” (such as in the phrase “strong AI” (Searle 1980)) fall within the scope of our account.

^{*}These authors contributed equally.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

2 Between Human Intelligence and Technology: AGI’s Dual Value-Laden Pedigrees

Building on STS scholarship on the values embedded in technology (Winner 1980), research communities like FAccT, AIES, and CHI have taken a deep interest in the political, social, and ethical values embedded both in AI tools and in the practices that surround AI (Fishman and Hancox-Li 2022; Dotan and Milli 2020; Birhane et al. 2022b; Scheurman, Hanna, and Denton 2021; Hutchinson et al. 2022; Bommasani 2022; Denton et al. 2020, 2021; Mathur, Lustig, and Kaziunas 2022; Shilton 2018; Brousard et al. 2019; Green 2021; Blodgett et al. 2020; Viljoen 2021; Abebe et al. 2020; Birhane and Guest 2021; Blili-Hamelin and Hancox-Li 2023; Costanza-Chock 2020). AGI and human-level AI concern not only existing technologies, but also the technologies that many AI builders, researchers, and organizations dream of building in the future. When companies describe their official “long term aim [as] to solve intelligence, developing more general and capable problem-solving systems, known as artificial general intelligence (AGI)” (DeepMind 2022), or when thought leaders claim that “[m]itigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war” (Center for AI Safety 2024), they engage in “the process of negotiating between competing perspectives, values, and goals” (Green 2021).

We interrogate the political, social, and ethical questions undergirding AGI discourse by paying attention to how different accounts embody competing visions for the future of technology. We don’t see value-laden assumptions as a flaw. We see them as choices that admit legitimate disagreement through competing values. Our account of alternative paths forward (Section 4) strives to be reflective and explicit about the political and social assumptions of the visions we call for—such as democracy, epistemic justice, contextualism, and participation. Throughout, we also embrace the value of reflectiveness (Boyarskaya, Olteanu, and Crawford 2020; Prunkl et al. 2021) about political, social, and ethical assumptions. We believe that making value assumptions explicit to ourselves and to others often makes for better individual and collective decisions—we might say, for more intelligent “experiments in living” (Mill 2003; Anderson 1991). Our aim in emphasizing reflectiveness is to point to more legitimate methods of determining the values that we want in AI. When large corporations are the only entities championing the narrative that their visions of AGI will transform everyone’s lives, what and whose values ultimately win out becomes a question of power divorced from legitimacy (Burrell and Metcalf 2024).

In the rest of this section, we examine a second and more often overlooked root of AGI’s value-laden character: *intelligence*. Debates about defining and measuring human intelligence are crucial to anticipating the value-laden aspects of AGI for at least three reasons. Firstly, many of the reasons for why definitions of human intelligence are value-laden carry over to the case of AI. Secondly, some attempts at defining AGI use concepts or methods from human intel-

ligence research. Thirdly, the question of whether machines match or surpass human intelligence faces the challenge of specifying what counts as human intelligence and of evaluating human intelligence.

2.1 Intelligence Is Value-Laden Because It Is a Thick Concept

Intelligence is what philosophers would call a *thick evaluative concept*: it includes both descriptive and normative elements (Anderson 2002; Kirchin 2013, 2017). It contains descriptive elements about what empirical phenomena fall under the concept of intelligent behavior. It also contains a normative element. Evaluating intelligence inevitably involves assessing the desirability of specific behaviors. Previous research (Anderson 2002; Alexandrova and Fabian 2022; Blili-Hamelin and Hancox-Li 2023; Cave 2020) has argued that when we try to define or measure thick concepts, we inevitably embed ethical values in the design decisions we make about the boundaries of the concepts and the measurement methodology. Other examples of thick concepts that have value-laden definitions are health and well-being (Alexandrova and Fabian 2022). The fact that intelligence is value-laden is at the root of debates over its definition and validity.

The validity of the concept of general intelligence in humans has been extensively criticized. One of the principle targets of these critiques is Spearman’s *g*: a hidden, not directly observable factor expressing “shared variance across a set of intercorrelating cognitive tasks” (Warne and Burningham 2019). Here we recap a few of the main critiques to provide some background for our later points about how some definitions of AGI make moves that are subject to similar critiques.

History of Ableism and Racism General intelligence metrics for humans have been criticized as ableist and racist (Anderson 2002; Summerfield 2023). Anti-racist critiques of human intelligence tests cite the influence of teacher expectations and stereotype threat, which affect test-takers of different demographics differently (Osborne 2001). Disability advocates have argued that conventional standardized tests for intelligence are ableist because they do not consider how disabled individuals can perform within a broader social context (Pellicano and den Houting 2022). In general, these critiques of *g* draw on the fact that human experiences and behaviors are diverse and contextual in ways that are hard to capture in standardized tests (Legault, Bourdon, and Poirier 2021).

Weaknesses of Factor Analysis The use of factor analysis to discover *g* as a causal factor has been extensively critiqued (Glymour 1998; Johnson 2016). These critics argue that factor analysis, on its own, does not rule out alternative causal structures that could explain the same empirical results. Factor analysis can show that a common cause is a possible explanation of observed statistical patterns. But those same patterns could be explained by many other causal hypotheses.

We mention this critique of factor analysis not just to situate the epistemic status of human intelligence metrics, but

also as background for our later discussion (Section 3.2) of attempts in AI to find the equivalent of g in AI systems. These attempts are subject to the same critique that factor analysis does not rule out alternative causal hypotheses.

One possible response to this critique of factor analysis is to treat g in a deflationary way: to say that it is not explanatory but is merely a shorthand for the correlation between performances on multiple cognitive tests. In this way, g could serve as a shorthand for *communicating* performance on a group of tests, without committing to it being a real causal factor. This is analogous to formulations of AGI that strive for a deflationary approach, as we describe in Section 3.1.

Circularity Another critique of psychometrics work purporting to quantify human intelligence is that it cannot be validated without reference to intelligence tests, making such measurements circular (Summerfield 2023; Boring 1923; Richardson 2017; Popper 2013). As Boring (1923) put it over 100 years ago, “[I]ntelligence is what the tests test”. Boring (1923)’s worry is not simply that intelligence tests need to be appealed to in validating intelligence tests. Rather, it is that the theoretical construct they purport to measure is itself motivated by setting aside the broad range of connotations of intelligence in favor of focusing on whatever can be measured through intelligence tests (Boring 1923). This critique points to how psychometrics research is insufficiently motivated and under-theorized. Warne and Burningham (2019) is an example of recent work that doubles down on disconnecting research on g from any attempt at anchoring psychometrics research in a substantive theoretical conception of intelligence. As with the critique of factor analysis (2.1), deflationary approaches that avoid ascribing an explanatory role to g may be a potential response to this worry.¹

2.2 Relevance to AGI

The critiques of g and the recognition of intelligence as a thick concept provide the background for why the definition and measurement of human intelligence is a political and social question. For similar reasons, defining and measuring AGI is also a political and social question. AGI is also a thick concept involving both descriptive (what tasks or abilities fall under its definition) and normative (what counts as *good* machine behavior) criteria. And, as we argue next, current definitions of AGI embed political and social values in ways that are often not explicitly acknowledged by their authors. In the next section, we identify the design choices that formulations of AGI make, emphasizing their contingency and the types of values that each choice embeds.

3 The Motley Choices of AGI Discourse

In this section, we look at some of the typical moves made when defining AGI. Our central contention is that these definitions always depend on assumptions about what is valuable to some group of people or what goals are

¹For an overview of recent approaches to motivating psychometrics research, see Deary (2012).

worthwhile—whether implicitly or explicitly. This is expected once we understand that the topic of AGI and human-level AI sits at the intersection of two fundamentally value-laden questions. The topic of intelligence is fundamentally value-laden due to its nature as a thick concept (see Section 2). Moreover, the topic of what technologies and tools are *worth building* is itself fundamentally value-laden (Costanza-Chock 2020; Birhane et al. 2022b; Dotan and Milli 2020). No conception of AGI can escape the political, social, and ethical priorities that *do* and *should* shape any answer to the question of what is worth building, by whom, and for what purposes. Throughout this section, we highlight value-laden aspects of AGI discourse that pertain both to its character as a technology and to the nature of the task of defining intelligence.

Some accounts of AGI attempt to adopt a kind of neutral stance. They may *start* by declaring that they intend to have a deflationary, value-neutral take on AGI (similar to the deflationary view on g we sketched in Section 2.1), but even these ultimately end up going beyond a purely deflationary view (see Section 3.3). We start by sketching the deflationary view (Section 3.1), then proceed to look at some of the major AGI definitions out there and identify dimensions on which different definitions make different value-laden assumptions (Section 3.2).

3.1 Purportedly Deflationary Accounts of AGI

Some definitions of AGI start out with seemingly neutral goals, such as creating a shared language or common standards for measuring AGI, without (at least at first) specifying any normative goals beyond that. This can be viewed as similar to the deflationary view of g described in Section 2.1: where g is simply a correlational factor that you can measure from a battery of tests and does not necessarily represent any underlying causal or explanatory factor. One can adopt standards for the sake of being able to consistently compare systems against one another within those standards without assuming that those standards are the only “real” or “objective” standards.

An example of the deflationary view is represented (at least initially) in “Levels of AGI” (Morris et al. 2023). At the start of the paper, the authors claim that they are seeking a “common language to compare models, assess risks, and measure progress along the path to AGI.” They think that having a “clear, operational definition of AGI” is the way to do this.

At first, their goal seems to be purely deflationary—they want some shared, operationalizable terminology so that we know when we are referring to similar phenomena. Morris et al. (2023) are inspired by the Levels of Automation framework (SAE International 2021) used to grade self-driving car technology—a framework that was also allegedly adopted to enhance communication.² Their deflationary approach is further underlined by a willingness to change the definition of AGI in the future if we gain the ability to operationalize things that are currently not operationalizable. For ex-

²However, see Section 3.3 for a discussion of how this framework also drifted from its initial deflationary goals.

ample, they currently want to define AGI in terms of “capabilities rather than processes” because we currently have little insight into the underlying mechanisms of many AI systems. They allow that research into mechanistic interpretability might eventually lead to operationalizable definitions involving processes, which then “may be relevant to future definitions of AGI” (Morris et al. 2023).

Another AGI account that is on the deflationary spectrum is that of Adams et al. (2012), who have “a pragmatic goal for measuring progress toward its attainment”. The authors acknowledge that “The heterogeneity of general intelligence in humans makes it practically impossible to develop a comprehensive, fine-grained measurement system for AGI”. However, in order to measure progress, they would still like “a common framework for collaboration and comparison of results”. This conception is similar to that of (Morris et al. 2023) in treating the AGI framework as a common language for communication purposes. To the extent that Adams et al. (2012) adopt requirements for “general cognitive architectures”, they emphasize that these requirements are not final and are “simply . . . a convenient point of departure for discussion and collaboration”.

We’ll see in Section 3.3 that these deflationary accounts depart from their initial declared intentions and end up including some value-laden choices. But first, we’ll take a tour of some non-deflationary accounts of AGI.

3.2 Value-Laden Choices in AGI Definitions: A Taxonomy

Having outlined deflationary views on AGI, we now identify some value-laden choices that non-deflationary views make in their conception of AGI or human-level AI. We organize these views by the dimensions on which these choices are made. Many of these choices are better understood as differences in focus and emphasis, rather than as mutually exclusive assumptions about the correct conception of AGI.

Economic Value Some definitions of AGI are explicit about the values that they are centering. For example, OpenAI’s charter claims that “OpenAI’s mission is to ensure that artificial general intelligence (AGI)—by which we mean highly autonomous systems that outperform humans at most economically valuable work—benefits all of humanity” (OpenAI 2018). This definition of AGI is meant to track the metric of *performing economically valuable work*. As others have pointed out (Morris et al. 2023), this assumes that other types of work are less valuable—a normative assumption.

OpenAI’s definition is perhaps the most explicit about its values. But we can find other definitions that embed similar values. Suleyman and Bhaskar (2023), for example, define “Artificial Capable Intelligence” by its ability to pass what they call the “Modern Turing Test”: the capability to turn a starting pot of \$100,000 of capital into \$1,000,000 over several months. Similarly, Nilsson (2005) proposes an “employment test” to replace the Turing Test: “To pass the employment test, AI programs must be able to perform the jobs ordinarily performed by humans. Progress toward human-level AI could then be measured by the fraction of these jobs

that can be acceptably performed by machines.”

In contrast, Morris et al. (2023) reject definitions based on economic value alone. They think that benchmark tasks should align with what people value beyond economic value, including metrics that are harder to automate or quantify. Here, in tension with their purported deflationary orientation (see Section 3.1), they explicitly argue for having components of AGI that are less operationalizable (because they are harder to quantify). The reason they are doing so is value-laden: because they think that there are values beyond economic value that are worthwhile to aim for.

Embodiment The choice of whether to include the ability to carry out tasks in the physical world is another dimension on which AGI frameworks differ. Morris et al. (2023) make an unexplained choice to exclude embodied tasks from their framework. This exclusion is particularly mysterious because physical tasks are, on the face of it, no harder to operationalize than non-physical tasks.

In contrast, criteria like Wozniak’s coffee test require embodied AI systems. This test, first proposed by Apple founder Steve Wozniak (Fast Company 2010), tasks the machine with going into an “average” American home and making coffee in it. It has since been included in definitions of AGI (Goertzel, Iklé, and Wigmore 2012; Marcus 2022). AI critics like Weizenbaum (1976) have also cited the field’s relative inattention to embodiment as a major shortcoming.

The decision to include or exclude embodiment in the definition of AGI is value-laden, as it incorporates assumptions about which tasks are considered valuable. Privileging the mental over the physical is a normative choice. It also carries normative consequences, such as influencing the types of AI systems likely to be developed and their potential impacts on society.

Human-Like Processes Versus Outcomes Another divisive choice in conceptualizing AGI or human-level AI is whether to insist on mechanisms or processes that mimic human cognitive processes. For example, Morris et al. (2023) want to prioritize outcomes (“what an AGI can accomplish”) over human-like processes because outcomes are more operationalizable. Excluding “processes” from the definition means that criteria based on the following are excluded: consciousness (Butlin et al. 2023; Lenharo 2024; Searle 1980; Summerfield 2023; Smart 2015), sentience (Schwitzgebel and Garza 2015), and abilities modeled on human children (Turing 2012; Nilsson 2005; Gopnik 2019, 2015; Summerfield 2023).

This choice of excluding anthropomorphic cognitive processes goes beyond purely conceptual questions about how to best define AGI.³ It also influences the types of research

³Here is a possible epistemic argument for taking a stance on the exclusion of anthropomorphic processes from the concept of AGI. The question of whether or not implementing human-like processes is needed to achieve human-like outcomes is an empirical question. Settling open-ended empirical questions through definitions is bad epistemic practice. Definitions of AGI should, therefore, be agnostic on empirical questions such as the importance of human-like processes to achieve different human-like outcomes. We see this as amounting to providing a principled reason for *not*

projects that attract AI investment, which has ethical consequences in terms of what the AI we build (and its attendant social consequences) looks like. Lenharo (2024) interviews a researcher claiming that “to his knowledge, there was not a single grant offer in 2023 to study the topic” of consciousness in AI.

Rich Sutton’s “The Bitter Lesson” articulates this practical concern about maximizing return on investment (Sutton 2019). Sutton frames this as a choice between focusing on implementing human-like processes versus focusing on “general methods that leverage computation”. He frames the choices as practically exclusive in the sense that “[t]ime spent on one is time not spent on the other” Sutton (2019). The bitter lesson is that building human-like processes into AI may be psychologically satisfying to researchers, but tends to be much less effective than focusing on “general purpose method[s]” that scale with computation.

However, it’s not clear that the practical argument for ignoring the processes underlying intelligence is *correct*. Summerfield (2023), for instance, makes the case that in drawing lessons from “natural general intelligence” (i.e. human cognition and intelligence), work on AGI stands to yield more “useful” and “effective” systems, and to better leverage the “tight intellectual synergy between AI research, cognitive science, and neuroscience”.

In short, different practitioners make different choices about what they consider to be “practical” for achieving AGI. Choices like these about effectiveness and how to best allocate limited resources are inherently value-laden—raising questions like *for what goals, given whose beliefs and preferences, given what conditions?* For example, Summerfield (2023) does not articulate what he means by “useful” or “effective”, let alone for whom it is useful or effective. Similarly, when Morris et al. (2023) argue for having shared, operationalizable (and thus outcome-based) standards for AGI on the grounds that the standards will be “useful”, they do not articulate whom it will be useful for, or consider whether it may be less than useful to some. Critics of current trends in AI have pointed to how its harms and benefits are distributed unevenly in ways that are correlated with existing power structures (Green 2021; Abebe et al. 2020; Blodgett et al. 2020; Birhane and Guest 2021; Costanza-Chock 2020). Given this context about how AI is currently used, implicit assumptions that “useful” means “useful for everyone” need more justification.

Generality A point of widespread agreement among accounts that specifically favor the term “AGI” over “human-level intelligence” is the importance of *generality* (Goertzel 2014; Morris et al. 2023; Agüera y Arcas and Norvig 2023; Summerfield 2023). Summerfield (2023) summarizes this as the view that a “generally intelligent individual is polymathic—good at everything.”⁴ The recent populariza-

answering the question of whether AGI needs to involve anthropomorphic cognitive processes. Discourse on AGI often goes beyond this agnostic stance for practical reasons.

⁴Summerfield (2023) emphasizes the importance of the kind of flexibility that humans display in ordinary everyday activities—“whether performing daily rituals of basic hygiene, navigating

tion of the term “AGI” is often traced to the mid-2000s interest in contrasting “narrow AI” capable of “specific ‘intelligent’ behaviors in specific contexts” with systems that would be able to “self-adapt to changes in their goals or circumstances”, “generalize knowledge from one goal or context to others”, and so on (Goertzel 2014).⁵

Although widespread, the strong emphasis on “generality, adaptability and flexibility” (Goertzel 2014) in machine intelligence is not universal among accounts of human-level machine intelligence. For instance, Bostrom (2014) departs from this norm by denying that achieving human-level intelligence or superintelligence requires “general” intelligence at all. Bostrom sees this as part of his strong rejection of the requirement for human-like cognitive processes. He instead remains agnostic about the structure of abilities needed to match or outclass the instrumental performance of human intelligence.

We argued earlier (Section 3.2) that assumptions about whether AGI should include human-like processes or not is value-laden because it is based on notions of “usefulness” or “effectiveness” that presuppose whom and what the assumptions are useful for. The same line of thought applies to arguments about whether the generality of abilities is important. Assumptions about what practical outcomes are desirable are needed to make the case for or against the importance of generality. Bostrom’s line of argument presumes that replicating the *instrumental* performance of human intelligence is the goal—but there are other possible goals for AI. Similarly, others’ assumption that being polymathic is essential to AGI glosses over two points: whether this is a desirable goal for AI, and the real diversity of human abilities (not all humans are similarly polymathic). Here we see an analogy with critiques of *g* on the basis that it ignores human diversity (see Section 2.1).

Individualism Conceptions of AGI also differ on whether they conceive of intelligence as a property of individuals—such as isolated humans or systems. One aspect of individualism about AGI is *ontological*, having to do with whether

the streets of their neighbourhood, negotiating the local market, deftly parenting their children, or judiciously managing household finances”—rather than in the kinds of abilities associated with “chess grandmasters” or “Nobel laureates”. This sense of generality is quite different from what gets emphasized by *g*. Rather than centering a putative variable that correlates with *improved performance* across all cognitive tasks, Summerfield emphasizes the multiple complex facets of the kind of cognitive flexibility that humans display.

⁵Goertzel (2014) traces the genealogy of the term “AGI” as follows. “The brief history of the term “Artificial General Intelligence” is as follows. In 2002, Cassio Pennachin and I were editing a book on approaches to powerful AI, with broad capabilities at the human level and beyond, and we were struggling for a title. I emailed a number of colleagues asking for suggestions. My former colleague Shane Legg came up with “Artificial General Intelligence,” which Cassio and I liked, and adopted for the title of our edited book [Goertzel et al. (2007)]. The term began to spread further when it was used in the context of the AGI conference series. A few years later, someone brought to my attention that a researcher named Mark Gubrud had used the term in a 1997 article on the future of technology and associated risks [Gubrud (1997)].”

the entity ascribed “intelligence” is an individual. The psychometrics project of measuring and comparing the intelligence and cognitive skills of humans takes individuals as a key unit of analysis. AI evaluation practices like benchmarking, as currently practiced, mostly treat individual models as bearers of the properties they measure, and definers of AGI often propose tests that are mostly or entirely tests on individual agents (Chollet 2019; Morris et al. 2023).

By contrast, accounts like Bostrom (2014) and Attard-Frost (2023) reject ontological individualism. Bostrom argues that when thinking about what it means for machines to match or outclass human intelligence, the relevant unit of comparison includes not just the intelligence of an individual human but also “the combined intellectual capability of all of humanity” at present (Bostrom 2014). Specifically, he argues that collectives—such as “firms, work teams, gossip networks, advocacy groups, academic communities, countries, even humankind as a whole” (Bostrom 2014)—can be understood as a mechanism for increasing intelligence. Bostrom also refers to collective intelligence as “collective intellectual problem-solving capacity”. This rejection of ontological individualism is related to the longer tradition of thinking about collective intelligence (Lévy 1994, 2010; Engelbart 1962; Baltzersen 2021; Suran, Pattanaik, and Draheim 2021; Araya and Marber 2023; Landemore and Elster 2012; Anderson 2006; Putnam 2011; Dewey 2011; Peters and Heraud 2015). Baltzersen (2021) proposes thinking of collective intelligence as “collective problem solving”, both in large and small groups of people.⁶ In Section 5, we come back to a strand of this tradition that frames *democracy* as a “precondition for the full application of intelligence to the solution of social problems” (Putnam 2011).

Ontological individualism has normative implications in shaping the “north star” goals of AI: is the imagined challenge building machines that rival the intelligence of isolated human individuals, or creating systems that replicate the collective intelligence of groups or entire societies? It is unclear, on the face of it, which goal is ethically preferable to the other (of course, not building either is also an option). Due to contemporary AI’s focus on individual intelligence, it is hard to envision how AI systems would be different if the focus was instead on collective intelligence—but it is plausible that that counterfactual would lead to very different AI systems, with correspondingly different impacts on society.

A second dimension of individualism in conceptions of AGI is methodological: concerning whether measurements of intelligence are carried out in environments where individuals are acting by themselves or where agents interact in a social environment with other agents. Most AI benchmarks are focused on individual agents carrying out tasks by themselves. But an alternative vision is possible (Attard-Frost 2023)—one that we describe in more detail in Section 4.1. These two different ways of measuring intelligence would lead to different prioritizations of resources for research and development, with attendant ethical implications. For exam-

⁶Lévy (2010) similarly defines collective intelligence as “the capacity of human collectives to engage in intellectual cooperation in order to create, innovate and invent”.

ple, if we measure intelligent agents’ performance in more social and interactive environments, would this also lead to AI systems having fewer unanticipated harmful effects when deployed into the “real world”? Arguably, many of the recent cases of unanticipated harmful effects from AI systems being deployed are due to these systems being evaluated pre-deployment on decontextualized tasks that ignore aspects of the social environment (Wolf, Miller, and Grodzinsky 2017; Ganguli et al. 2022; Saisubramanian, Zilberstein, and Kamar 2022).

Instrumentality Thesis A widespread assumption in discussions of AGI is what Russell (2019) calls the “standard model of intelligence”: “[Humans or] machines are intelligent to the extent that their actions can be expected to achieve their objectives.” Similarly, Legg and Hutter (2007) characterize intelligence as “measur[ing] an agent’s ability to achieve goals in a wide range of environments.” These accounts endorse what we call the *instrumentality thesis*: the view that intelligence is a means to whatever ends, final goals, or preferences an agent might have.⁷ These accounts conceive of being more or less intelligent as categorically different from being better or worse at determining what matters, what goals are worth pursuing, or what is good for its own sake. Bostrom (2003) illustrates this feature of intelligence with the example of a paperclip maximizer: an agent with the final goal of building a world with maximally many paperclips. Paperclips are sometimes useful tools, such as for organizing paper documents. But paperclip maximizing as a final goal is patently not worthwhile. The instrumentality thesis is part of a common tradition in AI: thinking of intelligent agents as maximizing the expected utility of actions for satisfying given objectives or preferences.

The assumption that intelligence is purely about instrumental rationality is another juncture where values come into play. An alternate conception of intelligence could consider the ability to determine what goals are worth pursuing as a form of intelligence. Agents with this ability would be able to form independent views about whether the goals they are assigned by other agents are worthwhile, and if not, how or whether to resist their assignments. In short, the view that intelligence encompasses only instrumental rationality lays the ground for AI to develop agents that can optimize the goals of their designers, but that have no ability to question those goals. We can imagine instead a very different future of AI: where we have intelligent agents that are able to dissent about final goals (see Section 5). Which future is more desirable is a question of values: do we want intelligent machines that are completely subservient, or do we want machines that can collaborate with humans in determining what

⁷Bostrom (2014) calls this the orthogonality thesis, which he defines as follows (Bostrom 2014): “Intelligence and final goals are orthogonal: more or less any level of intelligence could in principle be combined with more or less any final goal.” We find it more helpful to highlight the similarity of this conception of intelligence with the conceptions of instrumental rationality that come from traditions like utility theory, or embrace Hume’s adage that “Reason is, and ought only to be the slave of the passions” (Hume 1991; Schmitter 2021).

goals are worth pursuing?

Operationalizability/Measurability Another dimension on which accounts of AGI differ is whether to require operationalizability or measurability for their definition. At one end of the spectrum, we have Morris et al. (2023), who frame their project around operationalizability and explicitly reject criteria that are not operationalizable. However, other discussants include definitions that are not measurable. Bostrom (2014), for example, makes frequent reference to the collective intelligence of all of humanity at present in his framework—a concept that seems difficult, if not impossible, to operationalize.

Operationalizability may appear value-neutral at first glance. However, scholars have argued that simply having a common language and shared standards is in itself value-laden. Commensuration (the act of establishing common standards to measure things) allows large institutions to coordinate actions and automate decision-making, but less powerful groups often argue that some things cannot be measured in order to resist actions that can be justified through “rational choices” (Espeland and Stevens 1998). Social scientists have also recognized that commensuration has performative effects, creating a kind of path dependency: “as commensuration gets built into practical organizations of labor and resources, it becomes more taken for granted and more constitutive of what it measures” (Espeland and Stevens 1998). This type of path dependency has been recognized in various domains of commensuration, such as standardized grades of grain quality (Porter 1996) and university rankings (Fowles, Frederickson, and Koppell 2016). Previous work on ML benchmarking has also highlighted how benchmarking creates performativity and path dependency (Dehghani et al. 2021; Blili-Hamelin and Hancox-Li 2023).

Choice of Tasks and Benchmarks One cross-cutting dimension of difference across definitions of AGI is the tasks or benchmarks that they include in their criteria for AGI. Some of the previously mentioned dimensions intersect with this—for example, if embodiment is part of your AGI definition, you would include some embodied tasks in your definition. In contemporary AI research, with its plethora of tasks and benchmarks, researchers sometimes make choices about which tasks they emphasize when estimating progress toward AGI.

Two prominent AI researchers recently argued that “[t]he most important parts of AGI have already been achieved by the current generation of advanced AI large language models” (Agüera y Arcas and Norvig 2023). While these researchers define AGI roughly as instructable systems able to operate over a wide variety of topics, tasks, modalities and languages, the fact that the strengths of current LLMs are considered to already cover “the most important parts of AGI” implies that the things that LLMs cannot do are the less important parts of AGI. These LLM weaknesses include (as they acknowledge themselves) arithmetic and adhering to facts. However, they do not explain why these areas of reasoning are devalued relative to LLM strengths like translation or synthesizing code. Other researchers, in contrast,

believe that arithmetical and logical reasoning are necessary components of AGI (Marcus 2023). Agüera y Arcas and Norvig (2023)’s assertion that the most important aspects of AGI are achieved by current-day LLMs is value-laden. Choices of tasks and benchmarks are a key site of value-laden design choices (Blili-Hamelin and Hancox-Li 2023; Bommasani 2022). Deciding to focus on LLMs over deductive reasoning systems has ethical implications because these systems will change our societies in different ways.⁸

Importing *g* Into AI One family of AGI characterizations relies on directly importing the concept of *g* from human intelligence into AI. These papers assume that the hypothesis that *g* is an explanatory causal factor in human intelligence is true, then seek to discover it in AI systems. For example, Hernández-Orallo et al. (2021) try to find *g* by conducting factor analysis on the results of machine experiments. They motivate this by analogy to the supposedly explanatory role of *g* in human intelligence, which they do not question. This move has the following weaknesses.⁹

Firstly, the critiques of factor analysis as a methodology for finding a unique causal structure, as applied to the case of *g* in human intelligence (see Section 2.1), would also transfer over to the case of AI. Secondly, taking AI to have a directly analogous *g* factor means that all the value-laden assumptions around what human abilities count as important are imported over into the AGI case (see Section 2). For example, Chollet (2019), who also draws inspiration from *g*, takes for granted that there is a “space of tasks and domains that fit within the human experience” that we should use to measure intelligence—ignoring critiques of similar assumptions in the human intelligence literature (see Section 2.1).

3.3 Are Deflationary Views Truly Deflationary?

We’ve outlined dimensions on which accounts of AGI make value-laden choices that aren’t determined by purely epistemic criteria. We also saw that design choices made within the more “deflationary” views of AGI go beyond the apparently value-neutral goal of having a common language and shared standards. Here we discuss some more subtle ways in which these “deflationary” views incorporate values.

Morris et al. (2023) propose a common language in order to enable us to compare models, assess risks, and measure progress. These are different goals that a common language can potentially achieve. But it is not clear that each goal would be maximized by the *same* language. Should the shared standards that would be “best” for risk assessment necessarily be the same as those that would be best for measuring progress? To determine what shared standards would be best for risk assessment, we need to make some normative assumptions—for example, about which risks are more important. Our shared standards might differ depending on whether we think so-called “existential risk” is the biggest risk, in contrast to ongoing harms from AI.¹⁰ As an example,

⁸For another critique of the reliance on benchmarks in relation to AGI, see Summerfield (2023).

⁹For another critique, see also Russell (2019).

¹⁰For commentary on how policymakers are over-focused on existential risk, see (Hanna and Bender 2023) and (Milmo 2023).

later on in the paper, Morris et al. (2023) reject Suleyman and Bhaskar (2023)’s definition of AGI (see Section 3.2) because it might introduce “alignment risks”. It follows that minimizing alignment risks is at least an implicit desideratum for Morris et al. (2023)’s framework—but this is not stated upfront.

Looking at another deflationary view of AGI, Adams et al. (2012)’s way of framing the quest for AGI, as described in Section 3.1, assumes that their “pragmatic goal” is worth achieving—a topic on which people with different values might disagree.

Furthermore, there are insights to be drawn from Morris et al. (2023)’s analogy to the Society of Automotive Engineers’ (SAE) framework of levels of automation for automobiles (SAE International 2021). The intention of the analogy was to justify a similar move in the AI space, with the implied utility of a common language and shared standards. However, critical work on the SAE framework has argued that what appears to be merely a descriptive, technical definition with the SAE’s stated intentions of “simplifying communication” and “providing clarity” in fact has normative assumptions and implications (Hopkins and Schwanen 2021). Implications teased out by Hopkins and Schwanen include: promoting “homogeneity in mobility futures”, specifically a homogeneous vision of automation as the future of mobility; reproducing a “dominant discourse of an expert-led, technologically-centred vision of mobility futures”; and removing “obstacles and impediments to the successful (and timely) development of the automated vehicle niche.” Similarly, we may ask if shared standards for AGI promote homogeneity in AI futures and reproduce an expert-led dominant discourse. Certainly, the standards currently being developed are expert-led and lack input from people who are more likely to suffer negative impacts in AI—an issue that we discuss further in Section 4.3.

4 Towards Contextualized, Politically Legitimate, and Social Intelligence

We now outline alternative visions for values worth centering in imagining future forms of machine intelligence. These views embrace the value-laden nature of intelligence instead of side-stepping it. Broadly speaking, the following proposals take seriously the role of physical and social contexts in definitions of machine intelligence, in contrast to the views outlined in Section 3. We selected these views because they embody the values of contextualism, epistemic justice, inclusiveness, and democracy, which we consider vital for visions of the future of AI worth pursuing.

4.1 Contextual Intelligence

The importance of a contextual understanding of sociotechnical systems is well-recognized (Selbst et al. 2019; Weidinger et al. 2023; NIST 2023; Lazar and Nelson 2023; Shelby et al. 2023; Mohamed, Png, and Isaac 2020). Given the interlocking social and technical factors that shape the impact of AI systems, we need to consider “how a system is used, its interactions with other AI systems, who operates it, and the social context in which it is deployed”

(NIST 2023). Attard-Frost (2023)’s account of intelligence as “value-laden cognitive performance” brings a contextualist perspective to conceptualizing intelligence itself.

Their account centers the role of values in defining what counts as good cognitive performance, in addition to the role of “interdependencies between agents, their environments, and their measurers in collectively constructing and measuring context-specific performances of intelligent action” (Attard-Frost 2023). As they argue, contextually situated accounts of intelligence have a history in STS (Hayles 2017; Barad 2003), but also in AI (Weizenbaum 1976).

Attard-Frost’s account of intelligence contrasts with more individualistic views that measure intelligence “with reference to an extremely constrained and highly standardized set of cognitive activities performed by individuals, rather than with reference to situated activities performed in relation to other individuals and social environments” (Attard-Frost 2023). To use their lab analogy, Attard-Frost’s account is *in vivo* rather than *in vitro*—an account of how cognition is performed in social environments.¹¹ Another way of viewing the difference is that *in vivo* accounts have higher ecological validity—a property that even the less contextual accounts of AGI agree is desirable (Morris et al. 2023).

Attard-Frost’s account may have some practical drawbacks—perhaps it is harder to operationalize, or to compare systems that operate in different contexts. Thus, whether or not we choose to adopt something like their account is itself dependent on our values and goals: their goal of queering intelligence is itself a normative one, and it may conflict with other normative goals.

4.2 Inductive Risk and Social Values

Another helpful lens on how to define future forms of machine intelligence is that of inductive risk. In the values in science literature, inductive risk is often used to justify having different standards of evidence for accepting scientific claims. The idea is that when the potential social costs of accepting/denying a scientific claim are high, social and political values *should* influence standards of evidence (Douglas 2000; Steel 2010; Blili-Hamelin and Hancox-Li 2023).

We can adapt the argument from inductive risk to the case of defining AGI by reframing it in terms of standards of evidence for accepting definition(s) of AGI.¹² As Section 3 has illustrated, current AGI definitions make many choices that are unjustified by epistemic values, or are based on implicit assumptions about what is socially valuable. Adapting the argument from inductive risk to AGI, we would say: due to the significant social impacts that could arise from having definitions of intelligence that value certain social goals/tasks/beings over others, social and political values should influence how we define AGI.

¹¹See also Hutchins (1995) on the distinction between cognition “in the wild” and cognition under artificially controlled (laboratory) conditions.

¹²We frame it this way because it is unclear whether definitions of AGI are “scientific claims” as traditionally construed by philosophers of science. We thank Ravit Dotan for proposing this point.

4.3 Epistemic Justice for Defining Future Forms of AI

One aspect of many definitions of AGI that we have analyzed here is that they mostly come from actors who have relatively more power and influence over the future of AI. At the same time, there is a dearth of voices in the conversation from people who are more likely to be harmed by deployed AI systems. Advocates of AGI and human-level AI imagine this technology as impacting almost everyone. Taking that ambition seriously requires processes that give a meaningful say to the communities who would be impacted by the technology. Dreams of future technologies should come hand in hand with participatory, inclusive, and—as we argue below—politically legitimate decision-making processes.

A positive vision for defining future forms of AI (or, taking a step back, even deciding if it should be pursued) would learn from the participatory ML and epistemic justice literature. We would love to see a vision for future forms of machine intelligence that is constructed through participatory methods (Birhane et al. 2022a; Delgado et al. 2023; Young et al. 2024), while still being aware that the products of these methods have their limitations (Sloane et al. 2020). Needless to say, these participatory methods should take care to include the perspectives of not just those who design or fund AI, but those who will be impacted by it in other ways.

The epistemic justice lens also highlights discussions of epistemic justice in scientific fields that have had to make similar assessments of difficult-to-measure concepts. Epistemic justice is the idea that different groups of people have, due to power differentials, differing levels of credibility or differing contributions to the concepts that underpin our shared knowledge, and this is an injustice because it can render unequal distributions of consequences (e.g. misunderstanding what sexual harassment is has more impact on women than on men) or make it harder for less powerful groups to understand their own experiences (Fricker 2007; Schmidt 2019).

Alexandrova and Fabian (2022) argue that scientific fields studying thick concepts like intelligence, health, wellbeing, and sustainability have a distinctive need for participatory processes. For research that has implications on real-world communities—such as fields influencing policy and lawmaking—they argue that relying on the personal values of people studying the phenomena is not enough. Research on thick concepts that affect real-world communities, they argue, ought to seek *legitimacy* for its values. Political legitimacy concerns how states and institutions avoid coercion and oppression in their use of power over their own citizens, such as through justifying their power to those they have power over (Peter 2017). Alexandrova and Fabian (2022) propose an epistemic constraint on those fields, requiring the values of the field to be legitimated through a participatory “political process that includes all the stakeholders of this research.”

Similarly, Elabbar (2023) has extensively examined how values enter into Intergovernmental Panel on Climate Change (IPCC) reports. He considers these reports to involve difficult curatorial choices about “which truth-apt

claims and representations to display in a given space at all, under rationing pressure”. These curatorial choices are value-laden. For example, 89% of government reviews in the IPCC’s Summary for Policymakers are from developed nations, indicating one way that powerful nations have more input into the report. Crucially, the climate change case points to the ineffectiveness of interventions like transparency in design choices. Elabbar argues that being transparent about the role of values in justifying design choices fails to empower many stakeholders. This is because “developing emission figures presumes advanced knowledge of carbon accounting methods and various other forms of specialist expertise”, and “[i]n cases where the effect of value choices is subtle and complex, lay audiences will typically lack the capacity to cash out transparency in the form of genuine alternatives that accord with their values”. Thus, he rejects transparency as a solution to this type of epistemic injustice.

We think a similar situation applies to the problem of defining AGI. AGI definitions are currently being developed by people who, as technical experts designing AI systems, have relatively more power, and the effects of their design choices are complex. Transparency into this process would likely be insufficient to empower most stakeholders to imagine alternatives to these visions—which is not to say that it isn’t desirable. Participatory methods are thus a key part of the solution.

Notably, by suggesting more participatory methods, we do not mean that every issue surrounding future forms of machine intelligence should be subject to a population-wide referendum. As the IPCC case illustrates, specialist expertise is indispensable to making good decisions on these issues. Luckily, modern democracies have several methods of ensuring that citizens have a voice in state decisions through mechanisms other than referenda. For example, representative democracy, where citizens elect representatives who make decisions on behalf of them, mixes elements of expert-led decision making with citizen participation. Similarly, polities have also experimented with citizen assemblies on important issues (Devaney et al. 2020; Huang and Tu 2017). These are small forums where a selected set of “ordinary” citizens debate issues as a mechanism for influencing public opinion and political decision-makers.

5 Conclusion: Politically Legitimate Intelligence

Donna Haraway’s “A Manifesto for Cyborgs” argues that despite its origin in “racist, male-dominant capitalism” and “militarism”, the figure of the human-machine hybrid can be repurposed towards resisting oppressive social categories in imagining the future (Haraway 1987; Forlano and Glabau 2024). Imagining “altering bodily functions” to allow humans to survive in space (Clynes and Kline 1960) opens the door to Ashley Shew’s interrogation of how technology reinforces ableism: “Everyone in space will be disabled” (Shew 2018; Forlano and Glabau 2024). The scale, power, and scope—or perhaps lack of scope, as Gebru and Torres argue (Gebru and Torres 2024; Burrell and Metcalf 2024)—

of dreams of AGI should raise serious doubt about whether the concept is worth similarly subverting. Whether through subverting or discarding AGI, we nevertheless believe in the need for more contextual, participatory, and democratic approaches to imagining future forms of intelligence and machines.

To that end, we have taxonomized the discordant value-laden choices of AGI discourse. We surfaced alternative paths that provide more contextual, participatory, and epistemically just perspectives on imagining future forms of AI. We suggested that the question of what forms of machine intelligence are worth pursuing calls for a plurality of contested and value-laden perspectives. We now conclude by highlighting perspectives that provide positive rejoinders for placing dissent, deliberation, and political legitimacy at the center of conversations about intelligence and future technologies.

Democratic legitimacy and social conceptions of intelligence are connected. Birhane et al. (2022a) defend the need for participatory AI to be supplemented with approaches that yield “stronger forms of validation and legitimacy” through democratic governance. Though they do not refer to it, their position is resonant with an intellectual tradition that sees democratic institutions as embodying a distinctive form of *intelligence* (Anderson 2006, 2023; Putnam 2011; Landemore 2013; Festenstein 2019; Alexander and Kitcher 2021; Dewey 2011; Festenstein 2023; Anderson 1991; Mill 2003). This tradition asks the question: how can we arrive at empirically informed solutions to social problems that avoid subordinating people to others? The answer is to consider democracy as its own form of “social intelligence”—as a collective process of deliberation and reasoning that can lead to iterative improvement on solutions and worthwhile goals.

Building on John Dewey’s conception of “democracy as the use of social intelligence to solve problems of practical interest”, Anderson (2006) argues that democratic institutions embody solving social problems without resorting to oppression or coercion. She argues that democracies are uniquely suited to investigating solutions to *public interest* problems, problems that need to be solved through deliberation (“votes and talk”) rather than through procedures like markets. They do so by pairing sources of legitimacy like procedural fairness and universal inclusion (through mechanisms like law, rights, and voting), experimentation (“revising [...] decisions on the basis of experience with their consequence”), and dissent (Anderson 2006).

Public interest solutions to problems like deforestation and sustainability (Agarwal 2001) require institutions capable of correcting their shortcomings and unintended consequences.¹³ “Just as the solution to scientific problems is to

¹³Anderson illustrates these features through Agarwal (2001)’s study of the effects of gender exclusion in limiting the equity and the efficiency of community forestry groups (CFG) as sustainable solutions to deforestation in India and Nepal. Agarwal highlights a gendered division of labor that made women the primary users of local forests while imposing “formal and informal obstacles” to their participation in CFGs Anderson (2006)—such as the combination of rules that limit participation to 1 per household with patriarchal norms that favor men as representatives of households,

do more science, the cure for the ailments of democracy is more democracy” (Anderson 2006). Similarly, for Dewey, this process of appraisal, discussion, and judgment is what constitutes “organized intelligence”:

Of course, there are conflicting interests; otherwise there would be no social problems [...] The method of democracy—insofar as it is that of organized intelligence—is to bring these conflicts out into the open where their special claims can be seen and appraised, where they can be discussed and judged in the light of more inclusive interests than are represented by either of them separately. (Dewey 1987)¹⁴

Putnam (2011) summarizes this view: democracy “is the precondition for the full application of intelligence to the solution of social problems”.¹⁵ This lens puts pressure on the narrow range of problems, tasks, and processes that take center stage in discussions of would-be human-level AI and AGI. Intelligent solutions to social problems should not be framed as a matter of finding optimal means for satisfying fixed preferences. Rather, they require procedures that support universal inclusion in interrogation, deliberation, and dissent about what counts as the “common good” (Putnam 2011) and the “public interest” (Anderson 2006). They require procedures with political legitimacy (Peter 2017). On this view, future forms of AI can provide intelligent solutions to social problems only if they are also objects of dissent and deliberative contestation, rather than systems designed from the top-down by “experts” only. The project of imagining “worlds” (Costanza-Chock 2020) and future technologies worth building should be one of collective “experiments in living” (Mill 2003; Anderson 1991), in which all impacted rights-holders hold decision-makers accountable.

Acknowledgments

We are thankful to Ravit Dotan for extensive input on early versions of this paper (see also 4.2). This paper took shape in response to a practical problem we faced in an ongoing position paper project on AGI: the lack of homogeneity in conceptions of AGI. We felt that a paper dedicated to investigating definitions of AGI would help move the broader project forward. We are deeply grateful to our collaborators on the ongoing position paper project: (in alphabetical order) Chris Graziul, El Mahdi El Mhamdi, Hananel Hazan, Katherine Heller, Mariame Tighanimine, Margaret Mitchell, Olivia Guest, Roxana Daneshjou, Shiri Dori-Hacohen, Talia Ringer, and Todd Snider.

References

Abebe, R.; Barocas, S.; Kleinberg, J.; Levy, K.; Raghavan, M.; and Robinson, D. G. 2020. Roles for computing in social change. In *Proceedings of the 2020 Conference* and “CFG meeting times that coincide with women’s household tasks”.

¹⁴Quoted by Festenstein (2019).

¹⁵See Festenstein (2019) for an overview of critiques of this family of views, both as interpretations of Dewey and as conceptions of democracy.

- ence on Fairness, Accountability, and Transparency, 252–260. Barcelona Spain: ACM. ISBN 978-1-4503-6936-7.
- Adams, S. S.; Arel, I.; Bach, J.; Coop, R.; Furlan, R.; Goertzel, B.; Hall, J. S.; Samsonovich, A.; Scheutz, M.; Schlesinger, M.; Shapiro, S. C.; and Sowa, J. F. 2012. Mapping the Landscape of Human-Level Artificial General Intelligence. *AI Magazine*, 33(1): 25–41.
- Agarwal, B. 2001. Participatory Exclusions, Community Forestry, and Gender: An Analysis for South Asia and a Conceptual Framework. *World Development*, 29(10): 1623–1648.
- Agüera y Arcas, B.; and Norvig, P. 2023. Artificial General Intelligence Is Already Here.
- Alexander, N. R.; and Kitcher, P. 2021. Educating Democratic Character. *Moral Philosophy and Politics*, 8(1): 51–80. Publisher: De Gruyter.
- Alexandrova, A.; and Fabian, M. 2022. Democratising Measurement: or Why Thick Concepts Call for Coproduction. *European Journal for Philosophy of Science*, 12(1): 7.
- Anderson, E. 2002. Situated Knowledge and the Interplay of Value Judgments and Evidence in Scientific Inquiry. In Gärdénfors, P.; Woleński, J.; and Kijania-Placek, K., eds., *In the Scope of Logic, Methodology and Philosophy of Science*, 497–517. Dordrecht: Springer Netherlands. ISBN 978-90-481-6145-4 978-94-017-0475-5.
- Anderson, E. 2006. The Epistemology of Democracy. *Episteme*, 3(1-2): 8–22. Publisher: Cambridge University Press.
- Anderson, E. 2023. Dewey’s Moral Philosophy. In Zalta, E. N.; and Nodelman, U., eds., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring 2023 edition.
- Anderson, E. S. 1991. John Stuart Mill and Experiments in Living. *Ethics*, 102(1): 4–26. Publisher: University of Chicago Press.
- Araya, D.; and Marber, P., eds. 2023. *Augmented education in the global age: artificial intelligence and the future of learning and work*. New York: Routledge. ISBN 978-1-00-323076-2. OCLC: 1372600734.
- Attard-Frost, B. 2023. Queering intelligence: A theory of intelligence as performance and a critique of individual and artificial intelligence. In *Queer Reflections on AI*. Routledge. ISBN 978-1-00-335795-7. Num Pages: 17.
- Baltzersen, R. K. 2021. *Cultural-Historical Perspectives on Collective Intelligence: Patterns in Problem Solving and Innovation*. Cambridge University Press, 1 edition. ISBN 978-1-108-98136-1 978-1-108-83374-5 978-1-108-98675-5.
- Barad, K. 2003. Posthumanist Performativity: Toward an Understanding of How Matter Comes to Matter. *Signs: Journal of Women in Culture and Society*, 28(3): 801–831.
- Benjamin, R. 2024. *Imagination: a manifesto*. New York, NY: W.W. Norton & Company, first edition edition. ISBN 978-1-324-02097-4. OCLC: 1379265421.
- Birhane, A.; and Guest, O. 2021. Towards Decolonising Computational Sciences. *Kvinder, Køn & Forskning*.
- Birhane, A.; Isaac, W.; Prabhakaran, V.; Diaz, M.; Elish, M. C.; Gabriel, I.; and Mohamed, S. 2022a. Power to the People? Opportunities and Challenges for Participatory AI. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO ’22. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394772.
- Birhane, A.; Kalluri, P.; Card, D.; Agnew, W.; Dotan, R.; and Bao, M. 2022b. The Values Encoded in Machine Learning Research. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, 173–184. Seoul Republic of Korea: ACM. ISBN 978-1-4503-9352-2.
- Blili-Hamelin, B.; and Hancox-Li, L. 2023. Making Intelligence: Ethical Values in IQ and ML Benchmarks. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. June 12–15, 2023, Chicago, IL, USA.
- Blodgett, S. L.; Barocas, S.; Daumé Iii, H.; and Wallach, H. 2020. Language (Technology) is Power: A Critical Survey of “Bias” in NLP. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5454–5476. Online: Association for Computational Linguistics.
- Bommasani, R. 2022. Evaluation for Change. Publisher: arXiv Version Number: 1.
- Boring, E. G. 1923. Intelligence as the Tests Test It. *New Republic*.
- Bostrom, N. 2003. Ethical Issues in Advanced Artificial Intelligence.
- Bostrom, N. 2014. *Superintelligence: paths, dangers, strategies*. Oxford: Oxford University Press, first edition edition. ISBN 978-0-19-967811-2. OCLC: ocn881706835.
- Boyarskaya, M.; Olteanu, A.; and Crawford, K. 2020. Overcoming Failures of Imagination in AI Infused System Development and Deployment. ArXiv Version Number: 3.
- Broussard, M.; Diakopoulos, N.; Guzman, A. L.; Abebe, R.; Dupagne, M.; and Chuan, C.-H. 2019. Artificial Intelligence and Journalism. *Journalism & Mass Communication Quarterly*, 96(3): 673–695.
- Bubeck, S.; Chandrasekaran, V.; Eldan, R.; Gehrke, J.; Horvitz, E.; Kamar, E.; Lee, P.; Lee, Y. T.; Li, Y.; Lundberg, S.; Nori, H.; Palangi, H.; Ribeiro, M. T.; and Zhang, Y. 2023. Sparks of Artificial General Intelligence: Early experiments with GPT-4. ArXiv:2303.12712 [cs].
- Burrell, J.; and Metcalf, J. 2024. Introduction for the special issue of “Ideologies of AI and the consolidation of power”: Naming power. *First Monday*.
- Butlin, P.; Long, R.; Elmoznino, E.; Bengio, Y.; Birch, J.; Constant, A.; Deane, G.; Fleming, S. M.; Frith, C.; Ji, X.; Kanai, R.; Klein, C.; Lindsay, G.; Michel, M.; Mudrik, L.; Peters, M. A. K.; Schwitzgebel, E.; Simon, J.; and VanRullen, R. 2023. Consciousness in Artificial Intelligence: Insights from the Science of Consciousness. ArXiv:2308.08708 [cs, q-bio].
- Cave, S. 2020. The Problem with Intelligence: Its Value-Laden History and the Future of AI. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 29–35. New York NY USA: ACM. ISBN 978-1-4503-7110-0.

- Center for AI Safety. 2024. Statement on AI Risk.
- Chollet, F. 2019. On the Measure of Intelligence. ArXiv:1911.01547 [cs].
- Clynes, M. E.; and Kline, N. S. 1960. Cyborgs and space. *Astronautics*.
- Costanza-Chock, S. 2020. *Design justice: community-led practices to build the worlds we need*. Information policy. Cambridge, Massachusetts: The MIT Press. ISBN 978-0-262-04345-8.
- Deary, I. J. 2012. Intelligence. *Annual Review of Psychology*, 63(1): 453–482.
- DeepMind. 2022. About.
- Dehghani, M.; Tay, Y.; Gritsenko, A. A.; Zhao, Z.; Houlsby, N.; Diaz, F.; Metzler, D.; and Vinyals, O. 2021. The Benchmark Lottery. ArXiv:2107.07002 [cs].
- Delgado, F.; Yang, S.; Madaio, M.; and Yang, Q. 2023. The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO '23. New York, NY, USA: Association for Computing Machinery. ISBN 9798400703812.
- Denton, E.; Hanna, A.; Amironesei, R.; Smart, A.; and Nicole, H. 2021. On the genealogy of machine learning datasets: A critical history of ImageNet. *Big Data & Society*, 8(2).
- Denton, E.; Hanna, A.; Amironesei, R.; Smart, A.; Nicole, H.; and Scheuerman, M. K. 2020. Bringing the People Back In: Contesting Benchmark Machine Learning Datasets. ArXiv:2007.07399 [cs].
- Devaney, L.; Torney, D.; Brereton, P.; and Coleman, M. 2020. Ireland's Citizens' Assembly on Climate Change: Lessons for Deliberative Public Engagement and Communication. *Environmental Communication*, 14(2): 141–146.
- Dewey, J. 1987. Liberalism and Social Action. In *Later Works, vol. 11*. Southern Illinois University Press.
- Dewey, J. 2011. Creative Democracy: The Task before Us (Reprint from 1939). In *The pragmatism reader: from Peirce through the present*. Princeton, NJ Oxford: Princeton University Press. ISBN 978-0-691-13705-6 978-0-691-13706-3.
- Dotan, R.; and Milli, S. 2020. Value-laden disciplinary shifts in machine learning | Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. *FAT* '20: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*.
- Douglas, H. 2000. Inductive risk and values in science. *Philosophy of science*, 67(4): 559–579.
- Elabbar, A. 2023. The curatorial view of assessment and the ethics of scientific advice: Beyond decisional autonomy towards distributive epistemic justice. Last accessed: Jan 13, 2023.
- Engelbart, D. 1962. Augmenting Human Intellect: A Conceptual Framework. Technical report, Stanford Research Institute.
- Espeland, W. N.; and Stevens, M. L. 1998. Commensuration as a social process. *Annual review of sociology*, 24(1): 313–343.
- Fast Company. 2010. Wozniak: Could a Computer Make a Cup of Coffee? [Online; accessed 17-January-2023].
- Festenstein, M. 2019. Does Dewey Have an “epistemic argument” for Democracy? *Contemporary Pragmatism*, 16(2-3): 217–241.
- Festenstein, M. 2023. Dewey's Political Philosophy. In Zalta, E. N.; and Nodelman, U., eds., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2023 edition.
- Fishman, N.; and Hancox-Li, L. 2022. Should attention be all we need? The epistemic and ethical implications of unification in machine learning. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, 1516–1527. Seoul Republic of Korea: ACM. ISBN 978-1-4503-9352-2.
- Forlano, L.; and Glabau, D. 2024. *Cyborg*. The MIT Press essential knowledge series. Cambridge, Massachusetts ; London, England: The MIT Press. ISBN 978-0-262-37776-8 978-0-262-37777-5.
- Fowles, J.; Frederickson, H. G.; and Koppell, J. G. 2016. University rankings: Evidence and a conceptual framework. *Public Administration Review*, 76(5): 790–803.
- Fricker, M. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press. ISBN 0-19-823790-1.
- Future of Life Institute. 2023. Pause Giant AI Experiments: An Open Letter.
- Ganguli, D.; Hernandez, D.; Lovitt, L.; Askell, A.; Bai, Y.; Chen, A.; Conerly, T.; Dassarma, N.; Drain, D.; Elhage, N.; El Showk, S.; Fort, S.; Hatfield-Dodds, Z.; Henighan, T.; Johnston, S.; Jones, A.; Joseph, N.; Kernian, J.; Kravec, S.; Mann, B.; Nanda, N.; Ndousse, K.; Olsson, C.; Amodei, D.; Brown, T.; Kaplan, J.; McCandlish, S.; Olah, C.; Amodei, D.; and Clark, J. 2022. Predictability and Surprise in Large Generative Models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, 1747–1764. Association for Computing Machinery.
- Geburu, T.; and Torres, E. P. 2024. The TESCREAL bundle: Eugenics and the promise of utopia through artificial general intelligence. *First Monday*, 29(4).
- Glymour, C. 1998. What Went Wrong? Reflections on Science by Observation and *The Bell Curve*. *Philosophy of Science*, 65(1): 1–32.
- Goertzel, B. 2014. Artificial General Intelligence: Concept, State of the Art, and Future Prospects. *Journal of Artificial General Intelligence*, 5(1): 1–48.
- Goertzel, B.; Iklé, M.; and Wigmore, J. 2012. The architecture of human-like general intelligence. In *Theoretical foundations of artificial general intelligence*, 123–144. Springer.
- Goertzel, B.; Pennachin, C.; Gabbay, D. M.; Siekmann, J.; Bundy, A.; Carbonell, J. G.; Pinkal, M.; Uszkoreit, H.; Veloso, M.; Wahlster, W.; and Wooldridge, M. J., eds. 2007. *Artificial General Intelligence*. Cognitive Technologies. Berlin, Heidelberg: Springer. ISBN 978-3-540-23733-4 978-3-540-68677-4.

- Gopnik, A. 2015. What babies tell us about artificial intelligence.
- Gopnik, A. 2019. AIs Versus Four-Year-Olds. In Brockman, J., ed., *Possible minds: twenty-five ways of looking at AI*. New York: Penguin Press. ISBN 978-0-525-55799-9 978-0-525-55801-9.
- Green, B. 2021. Data Science as Political Action: Grounding Data Science in a Politics of Justice. *Journal of Social Computing*, 2(3): 249–265. ArXiv:1811.03435 [cs].
- Gubrud, M. A. 1997. Nanotechnology and International Security. In *Nanotechnology and International Security*.
- Hanna, A.; and Bender, E. M. 2023. AI Causes Real Harm. Let's Focus on That over the End-of-Humanity Hype. *Scientific American*. [Online; accessed 17-January-2023].
- Haraway, D. 1987. A manifesto for Cyborgs: Science, technology, and socialist feminism in the 1980s. *Australian Feminist Studies*, 2(4): 1–42.
- Hayles, N. K. 2017. *Unthought: the power of the cognitive nonconscious*. Chicago ; London: The University of Chicago Press. ISBN 978-0-226-44774-2 978-0-226-44788-9.
- Heaven, W. D. 2020. Artificial general intelligence: Are we close, and does it even make sense to try?
- Hernández-Orallo, J.; Loe, B. S.; Cheke, L.; Martínez-Plumed, F.; and Ó hÉigeartaigh, S. 2021. General intelligence disentangled via a generality metric for natural and artificial intelligence. *Scientific Reports*, 11(1): 22822. Number: 1 Publisher: Nature Publishing Group.
- Hopkins, D.; and Schwanen, T. 2021. Talking about automated vehicles: What do levels of automation do? *Technology in Society*, 64: 101488.
- Huang, T.-y.; and Tu, W. 2017. public policy processes and Citizen participation in taiwan. In *Public Administration in East Asia*, 517–532. Routledge.
- Hume, D. 1991. *A treatise of human nature*. Oxford: Clarendon Pr, 2. ed., 8. impr edition. ISBN 978-0-19-824588-9.
- Hutchins, E. 1995. *Cognition in the Wild*. MIT Press.
- Hutchinson, B.; Rostamzadeh, N.; Greer, C.; Heller, K.; and Prabhakaran, V. 2022. Evaluation Gaps in Machine Learning Practice. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, 1859–1876. Seoul: ACM. ISBN 978-1-4503-9352-2.
- Ingram, D. 2024. Meta is pouring money into the creation of human-level artificial intelligence, Zuckerberg says. *NBC News*.
- Johnson, K. 2016. Realism and Uncertainty of Unobservable Common Causes in Factor Analysis. *Noûs*, 50(2): 329–355.
- Kirchin, S., ed. 2013. *Thick concepts*. Mind Association occasional series. Oxford: Oxford University Press, first edition. ISBN 978-0-19-967234-9.
- Kirchin, S. 2017. *Thick evaluation*. Oxford: Oxford University Press, 1st edition. ISBN 978-0-19-880343-0.
- Landemore, H. 2013. *Democratic reason: politics, collective intelligence, and the rule of the many*. Princeton ; Oxford: Princeton University Press. ISBN 978-0-691-15565-4.
- Landemore, H.; and Elster, J., eds. 2012. *Collective Wisdom: Principles and Mechanisms*. Cambridge University Press, 1 edition. ISBN 978-0-511-84642-7 978-1-107-01033-8 978-1-107-63027-7.
- Lazar, S.; and Nelson, A. 2023. AI safety on whose terms? *Science*, 381(6654): 138–138.
- LeCun, Y. 2022. About the raging debate regarding the significance of recent progress in AI.
- Legault, M.; Bourdon, J.-N.; and Poirier, P. 2021. From neurodiversity to neurodivergence: the role of epistemic and cognitive marginalization. *Synthese*, 199(5-6): 12843–12868.
- Legg, S.; and Hutter, M. 2007. A Collection of Definitions of Intelligence. ArXiv:0706.3639 [cs].
- Lenharo, M. 2024. AI consciousness: scientists say we urgently need answers. *Nature*, 625(7994): 226–226.
- Lévy, P. 1994. *L'Intelligence collective. Pour une anthropologie du cyberspace*. Paris: La Découverte.
- Lévy, P. 2010. From social computing to reflexive collective intelligence: The IEML research program. *Information Sciences*, 180(1): 71–94.
- Marcus, G. 2022. Dear Elon Musk, here are five things you might want to consider about AGI. [Online; accessed 17-January-2023].
- Marcus, G. 2023. Reports of the birth of AGI are greatly exaggerated. [Online; accessed 16-January-2023].
- Mathur, V.; Lustig, C.; and Kaziunas, E. 2022. Disordering Datasets: Sociotechnical Mismatches in AI-Mediated Behavioral Health. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2): 1–33.
- McCarthy, J.; Minsky, M.; Rochester, N.; and Shannon, C. 1955. A Proposal for The Dartmouth Summer Research Project on Artificial Intelligence. Technical report, Dartmouth Workshop.
- McLean, S.; Read, G. J. M.; Thompson, J.; Baber, C.; Stanton, N. A.; and Salmon, P. M. 2023. The risks associated with Artificial General Intelligence: A systematic review. *Journal of Experimental & Theoretical Artificial Intelligence*, 35(5): 649–663. Publisher: Taylor & Francis .eprint: <https://doi.org/10.1080/0952813X.2021.1964003>.
- Meaker, M. 2023. Meet Aleph Alpha, Europe's Answer to OpenAI. *Wired*. Section: tags.
- Mill, J. S. 2003. *On Liberty*. Yale University Press.
- Milmo, D. 2023. AI doomsday warnings a distraction from the danger it already poses, warns expert. *The Guardian*. [Online; accessed 22-January-2023].
- Mitchell, M. 2024. Debates on the nature of artificial general intelligence. *Science*, 383(6689): eado7069.
- Mohamed, S.; Png, M.-T.; and Isaac, W. 2020. Decolonial AI: Decolonial Theory as Sociotechnical Foresight in Artificial Intelligence. *Philosophy & Technology*, 33(4): 659–684.
- Morris, M. R.; Sohl-dickstein, J.; Fiedel, N.; Warkentin, T.; Dafoe, A.; Faust, A.; Farabet, C.; and Legg, S. 2023. Levels of AGI: Operationalizing Progress on the Path to AGI. ArXiv:2311.02462 [cs].

- Naudé, W.; and Barten, O. 2023. Artificial General Intelligence: can we avoid the ultimate existential threat?
- Nilsson, N. J. 2005. Human-Level Artificial Intelligence? Be Serious! *AI Magazine*, 26(4): 68–68. Number: 4.
- NIST. 2023. AI Risk Management Framework: AI RMF (1.0). Technical Report NIST AI 100-1, National Institute of Standards and Technology, Gaithersburg, MD.
- Noy, I.; and Uher, T. 2022. Four New Horsemen of an Apocalypse? Solar Flares, Super-volcanoes, Pandemics, and Artificial Intelligence. *Economics of Disasters and Climate Change*, 6(2): 393–416.
- OpenAI. 2018. OpenAI Charter. Technical report, OpenAI.
- Osborne, J. W. 2001. Testing stereotype threat: Does anxiety explain race and sex differences in achievement? *Contemporary educational psychology*, 26(3): 291–310.
- Pellicano, E.; and den Houting, J. 2022. Annual Research Review: Shifting from ‘normal science’ to neurodiversity in autism science. *Journal of Child Psychology and Psychiatry, and Allied Disciplines*, 63(4): 381–396.
- Peter, F. 2017. Political Legitimacy. In Zalta, E. N.; and Nodelman, U., eds., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2017 edition.
- Peters, M. A.; and Heraud, R. 2015. Toward a political theory of social innovation: collective intelligence and the co-creation of social goods. Accepted: 2015-08-27T04:47:37Z.
- Popper, K. 2013. *Knowledge and the body-mind problem: In defence of interaction*. Routledge.
- Porter, T. M. 1996. *Trust in numbers: The pursuit of objectivity in science and public life*. Princeton University Press.
- Prunkl, C. E. A.; Ashurst, C.; Anderljung, M.; Webb, H.; Leike, J.; and Dafoe, A. 2021. Institutionalizing ethics in AI through broader impact requirements. *Nature Machine Intelligence*, 3(2): 104–110.
- Putnam, H. 2011. A Reconsideration of Deweyan Democracy (Reprint from 1989). In *The pragmatism reader: from Peirce through the present*. Princeton, NJ Oxford: Princeton University Press. ISBN 978-0-691-13705-6 978-0-691-13706-3. Reprint of (1989). A Reconsideration of Deweyan Democracy. *Southern California Law Review*, 63, 1671.
- Richardson, K. 2017. *Genes, brains, and human potential: The science and ideology of intelligence*. Columbia University Press.
- Russell, S. J. 2019. *Human compatible: artificial intelligence and the problem of control*. New York?: Viking. ISBN 978-0-525-55861-3.
- SAE International. 2021. J3016.202104: Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles.
- Saisubramanian, S.; Zilberstein, S.; and Kamar, E. 2022. Avoiding negative side effects due to incomplete knowledge of AI systems. *AI Magazine*, 42(4): 62–71.
- Scheuerman, M. K.; Hanna, A.; and Denton, E. 2021. Do datasets have politics? Disciplinary values in computer vision dataset development. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW): 1–37.
- Schmidt, K. C. 2019. *Epistemic Justice and Epistemic Participation*. Ph.D. thesis, Washington University in St. Louis.
- Schmitter, A. M. 2021. 17th and 18th Century Theories of Emotions. In Zalta, E. N., ed., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2021 edition.
- Schwitzgebel, E.; and Garza, M. 2015. A Defense of the Rights of Artificial Intelligences: Defense of the Rights of Artificial Intelligences. *Midwest Studies In Philosophy*, 39(1): 98–119.
- Searle, J. R. 1980. Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3): 417–424.
- Selbst, A. D.; Boyd, D.; Friedler, S. A.; Venkatasubramanian, S.; and Vertesi, J. 2019. Fairness and Abstraction in Sociotechnical Systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 59–68. Atlanta GA USA: ACM. ISBN 978-1-4503-6125-5.
- Shelby, R.; Rismani, S.; Henne, K.; Moon, A.; Rostamzadeh, N.; Nicholas, P.; Yilla, N.; Gallegos, J.; Smart, A.; Garcia, E.; and Virk, G. 2023. Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction. ArXiv:2210.05791 [cs].
- Shew, A. 2018. Disabled People in Space – Becoming Interplanetary.
- Shilton, K. 2018. Values and Ethics in Human-Computer Interaction. *Foundations and Trends® in Human-Computer Interaction*, 12(2): 107–171.
- Sloane, M.; Moss, E.; Awomolo, O.; and Forlano, L. 2020. Participation is not a Design Fix for Machine Learning. arXiv:2007.02423 [cs]. ArXiv: 2007.02423.
- Smart, A. 2015. *Beyond zero and one: machines, psychedelics, and consciousness*. New York: OR Books. ISBN 978-1-68219-006-7. OCLC: 930340933.
- Steel, D. 2010. Epistemic values and the argument from inductive risk. *Philosophy of science*, 77(1): 14–34.
- Suleyman, M.; and Bhaskar, M. 2023. *The Coming Wave*. New York: Crown. ISBN 978-0-593-59396-7.
- Summerfield, C. 2023. *Natural general intelligence: how understanding the brain can help us build AI*. Oxford New York, NY: Oxford University Press, first edition edition. ISBN 978-0-19-284388-3.
- Suran, S.; Pattanaik, V.; and Draheim, D. 2021. Frameworks for Collective Intelligence: A Systematic Literature Review. *ACM Computing Surveys*, 53(1): 1–36.
- Sutton, R. 2019. The bitter lesson.
- Toews, R. 2022. Reflecting On ‘Artificial General Intelligence’ And AI Sentience.
- Turing, A. M. 1950. Computing machinery and intelligence (1950). *The Essential Turing: the Ideas That Gave Birth to the Computer Age*, 433–464.
- Viljoen, S. 2021. A Relational Theory of Data Governance. *The Yale Law Journal*.
- Warne, R. T.; and Burningham, C. 2019. Spearman’s g found in 31 non-Western nations: Strong evidence that g is a universal phenomenon. *Psychological Bulletin*, 145(3): 237–272.

Weidinger, L.; Rauh, M.; Marchal, N.; Manzini, A.; Hendricks, L. A.; Mateos-Garcia, J.; Bergman, S.; Kay, J.; Griffin, C.; Bariach, B.; Gabriel, I.; Rieser, V.; and Isaac, W. 2023. Sociotechnical Safety Evaluation of Generative AI Systems. ArXiv:2310.11986 [cs].

Weizenbaum, J. 1976. *Computer power and human reason: from judgment to calculation*. San Francisco: Freeman. ISBN 978-0-7167-0464-5 978-0-7167-0463-8.

Winner, L. 1980. Do Artifacts Have Politics? *Daedalus*, 109(1): 121–136. Publisher: The MIT Press.

Wolf, M. J.; Miller, K.; and Grodzinsky, F. S. 2017. Why we should have seen that coming: comments on Microsoft’s “taylor” experiment,” and wider implications. *ACM SIGCAS Computers and Society*, 47(3): 54–64.

Young, M.; Ehsan, U.; Singh, R.; Tafesse, E.; Gilman, M.; Harrington, C.; and Metcalf, J. 2024. Participation versus scale: Tensions in the practical demands on participatory AI. *First Monday*.