
Making Intelligence: Ethics, IQ, and ML Benchmarks

Borhane Blili-Hamelin*
borhane.blilihamelin@gmail.com

Leif Hancox-Li*
leif@fastmail.com

Abstract

The ML community recognizes the importance of anticipating and mitigating the potential negative impacts of benchmark research. In this position paper, we argue that more attention needs to be paid to areas of ethical risk that lie at the technical and scientific core of ML benchmarks. We identify overlooked structural similarities between human IQ and ML benchmarks. Human intelligence and ML benchmarks share similarities in setting standards for describing, evaluating and comparing performance on tasks relevant to intelligence. This enables us to unlock lessons from feminist philosophy of science scholarship that need to be considered by the ML benchmark community. Finally, we outline practical recommendations for benchmark research ethics and ethics review.

1 Introduction

Since 2020, the NeurIPS conference has been systematizing its approach to taking ethical risks into account in submitting and reviewing papers. As part of this work, the conference made public new ethical review guidelines in 2021 [47]. In 2022, a Provisional Draft of the NeurIPS Code of Ethics was made public and open to public input on OpenReview [5], addressing topics like the uses of research and the treatment of human subjects. In this position paper, we argue that the ethical risks of ML benchmarks go beyond those topics. Similarly to IQ tests, ML benchmarks involve ethical risks that trouble the line between technical or scientific concerns (such as construct validity and generality), and ethical concerns (such as justice, respect for persons, autonomy, etc.)

Our paper makes two contributions. First, we provide a conceptual framework for examining the junctures at which seemingly purely technical or scientific decisions about ML benchmarks do and should be informed by ethical considerations. Second, our paper provides a conceptual framework for understanding an overlooked similarity between IQ and ML benchmarks: they both set standards for quantitative description, evaluation, and comparison of performance on tasks relevant to intelligence. Both involve the specification of standard tasks for comparison of performance. Moreover, both involve quantitative metrics of success on those tasks, sometimes followed by a step of weighting tasks to calculate overall rankings across different tasks [76, 70].² We contend that the ML benchmark community stands to benefit from paying close attention to this overlooked structural similarity. Specifically, we show how insights from feminist philosophy of science scholarship on IQ research helps anticipate *areas of ethical risk* within technical and scientific decisions about ML benchmarks.³ These are:

*Names listed alphabetically. Both authors have contributed equally to the paper.

²By contrast, typical comparisons between ML and human intelligence tend to focus on similarities and differences between human cognitive abilities and the abilities of ML systems [8, 48]. Likewise, attempts have been made at building datasets that enable measuring ML model performance on human IQ tests [45]. This question of similarities and differences between human and ML abilities will not be the focus of our paper.

³Our approach takes on the burden of arguing that each of the points we make about benchmarks apply because of features of ML benchmarks, not because benchmarks are similar to IQ. Rather, we use the IQ case

1. The ethical risks that come with task selection—selecting which tasks matter enough to be included in a benchmark. (Section 4.1).
2. The ethical risks that come with choosing standards of construct validity, and with prioritizing specific forms of validity over others. (Section 4.2)
3. The ethical risks that come with path dependence, where benchmarks and particular types of models reinforce each other in a positive feedback loop. (Section 4.3).

After arguing for these claims, we outline (Section 5) practical recommendations towards factoring these areas of risk in research ethics and ethics review for the NeurIPS Datasets and Benchmarks Track. We conclude (Section 6) with a reflection on what good science means for ML benchmarks.

2 Background

The NeurIPS community already acknowledges that ethics review for benchmarks papers should pay close attention to issues like the potential uses of ML benchmarks and the treatment of human subjects [47, 5]. These are risks that are shared by benchmark and non-benchmark ML research papers. In this paper, we are interested in ethical risks that set benchmarks apart from other areas of ML research. We argue that much greater attention needs to be paid to ethical risks within the scientific and technical core of ML benchmarks.

Following Raji et al. [58], we understand ML benchmarks to involve a combination of dataset(s) and metrics, where the metrics attempt to capture performance or ability on task(s). Benchmarks influence practices in the ML research community and in the design of practical ML applications, where benchmarks are used to compare the performance of ML methods and models on some task(s). Influential benchmarks shape research agendas on the specific ML tasks they center and determine the dominant paradigms of ML research and applications.⁴ This practice of adopting benchmarks for comparison of model performance, we maintain, opens the door to areas of ethical risk that require extra attention from the NeurIPS benchmark community. In our view, this is also the area where lessons from the case of *human intelligence research* are especially helpful.

The areas of ethical risk we examine place our work at the intersection of two areas in recent work on the practices of ML research. In its first year, the NeurIPS Datasets and Benchmark track established itself as the premiere venue for position papers on ML benchmarks and datasets: interrogating practices like the prioritization of benchmarks on tasks believed to be indicators of progress on general-purpose ability, the prioritization of internal validity over external validity, and SOTA-chasing on a handful of influential benchmarks [58, 43]. These three tendencies overlap with the areas of ethical risk we identify: around task-selection, standards of validity, and path dependence.

Our paper also intersects with a growing body of research interrogating the ways in which technical or scientific priorities in ML research turn out to be *value-laden*: that is, dependent on social, political, and/or ethical values [22, 6, 27, 62].⁵ To our knowledge, our account is the first to focus on value-laden aspects of ML benchmarks—understood as combinations of datasets and metrics.⁶ Having in mind the context of ethics review, we focus on *ethical values* in a broad sense: the considerations that we rely on in assessing what actions are more or less worth pursuing or what outcomes are more or less worth bringing about.

to highlight how difficulties in defining IQ benchmarks are due to structural features in the problem that also exist in ML. We then independently argue that these same structural features of ML benchmarks lead to similar difficulties as those encountered in IQ.

⁴See Dotan and Milli [22] for a helpful look at the values implicit in ImageNet’s influence over the rise of deep learning as a dominant paradigm in ML research.

⁵For instance, Birhane et al. [6] argue that the most mentioned priorities in highly cited NeurIPS and ICML only appear purely scientific or technical on a superficial analysis. They argue that upon closer scrutiny, values like “performance, generalization, building on past work, quantitative evidence, efficiency, and novelty” tend to support the centralization of power in ML research.

⁶Recent work like [62] has started identifying values in datasets.

3 Values in Human Intelligence Research

What roles do ethical values play in measuring human intelligence? In this section, we examine this question with an eye to identifying helpful lessons for the case of ML benchmarks 4. This section also serves as an introduction to key philosophy of science concepts about the place of ethical values in scientific research: thick concepts in social science, the uses of research, and the argument from inductive risk.

Our perspective is deeply informed by feminist philosophy of science scholarship on the value neutrality of science. (See Appendix A.1 for more context.) We especially draw on Elizabeth Anderson’s investigation of the respects in which antiracist, feminist research on IQ challenges strong conceptions of the value-neutrality of science, while—as we see in Section 6 and A.1—preserving an important place for a limited form of value neutrality.⁷ This section reconsiders and bolsters Anderson’s key arguments, in order to later show how they transfer over to AI.

In the context of intelligence measurement, “intelligence” is used in a gradable sense, as picking out something that comes in degrees. Measurement of human intelligence is concerned with abilities that humans can have more or less of [60, 10]. “Intelligence” sometimes gets used in a categorical sense, as picking out a difference in kind: e.g. “are there intelligent life forms that don’t originate on our planet”? Or in the phrase, “machine learning is a subtype of artificial intelligence”. In this paper we focus on the gradable sense of intelligence.

Intelligence (in the gradable sense) is a thick evaluative concept [3, 10]. Thick concepts blur the line between evaluation and description. They convey both content about how we evaluate the world (approval or disapproval, praise or blame, success or failure, etc.) and about features of the world that seem independent of our evaluations [38, 39]. For instance, calling someone “open-minded” describes empirical features of how they tend to act across situations—and perhaps about their personality [69]. But open-mindedness is also a term of *epistemic praise*: expressing *values, standards, and/or ideals* having to do with how we should form and justify beliefs, how we should respond to evidence and/or uncertainty, how and when it is appropriate or praiseworthy to change one’s mind—say, when presented with new evidence.

Intelligence is a thick concept, in the sense that it is a term of praise that also purports to describe empirical facts about people’s abilities [3, 10]. Recognizing this feature of the concept of intelligence is helpful to identifying otherwise overlooked areas where ethical values and risks play a role in intelligence research. In Section 4, we will also see how it helps conceptualise the structural parallels between IQ and ML benchmarks.

We take the value neutrality of science to be at issue when ethical, social, or political values are used to make scientific decisions (see A.1). Many scientists recognize that the value of particular research projects or even areas of study can sometimes be outweighed by ethical values—this is why we have regulations about experiments involving humans and animals. However, in the case of measuring human and machine intelligence, values are embedded in a deeper way: in how we define the measures themselves (3.1), in our selection of standards of validity (3.2), and through path dependence (3.3).

3.1 Values determine what counts as intelligence

Values enter in defining the boundaries of the objects of intelligence research. Intelligence research is concerned with abilities. But what are their boundaries? Where do abilities begin and end? What marks an ability as relevant or irrelevant to intelligence? Whose abilities is intelligence research about?

Consider questions about the relationship between the objects of intelligence research and cultural boundaries. Are the findings of intelligence research culturally specific? Warne and Burningham [73] argue that this concern is heightened by the fact that definitions of intelligence are not only variable across cultures: even within specific cultures, there is a lack of expert consensus not only on what generally falls under the term “intelligence”, but also on the specific abilities that matter to intelligence.

⁷For an account of general parallels between the value-laden history of IQ research and AI, see Cave [10].

Following Anderson, we consider boundary problems to come hand in hand with the fact that intelligence is a thick evaluative concept [3]. There are cross-cultural variations and there is lack of expert consensus on what falls under “intelligence” because ethical values and interests play a central role in determining the what phenomena fall under the concept and the theoretical content of the concept.

This problem is shared with research on topics like *health* [3] and *well-being* [2]. Social science disciplines that inherit their subject matter from *thick concepts* usually face problems with separating the selection of the boundaries of their topic from *ethical* values.

As Alexandrova and Fabian argue, a common strategy towards sidestepping boundary problems with thick concepts in social sciences is trying to convert the thick concepts into technical terms [2]. This strategy remains recently favored by some human intelligence researchers. For the sake of securing cross-culturally invariable boundaries for the object of intelligence research, Warne and Burningham [73] propose that researchers should focus on an object whose boundaries are simply a matter of “statistical observation”: Spearman’s *g* [73].

What can be directly observed in cognitive tests is performance on very specific tasks [65, 9, 14, 16]. For instance, one test contains up to 15 different tasks: word similarity, vocabulary, visual puzzles, symbol search, digit span, comprehension, etc. [76, 16].

Early on, the field struggled with empirically studying mental abilities beyond performance on very specific tasks. Spearman is credited with realizing that these problems can be sidestepped through a technical procedure: statistically estimating whatever hidden factor reliably co-varies with observable performance on those tasks [65, 30, 9, 14, 73]. Spearman’s *g* is that hidden factor expressing “shared variance across a set of intercorrelating cognitive tasks” [73].⁸

Is reliance on *g* enough to eliminate reliance on ethical, social, and political values in determining the boundaries and objects of human intelligence research?

3.1.1 Quantitative definitions do not solve the boundary problem

First, the strategy of relying on *g* at best helps with only one of the many phenomena that interest intelligence researchers. A popular taxonomy in the field, introduced by Carroll [9], distinguishes three levels at which variations in performance on specific tasks occurs [9, 15]. Differences in performance on one mental task can correlate with: (1) variations in general performance on all mental tasks (the level of Spearman’s *g*); (2) variations in performance on a specific family of mental tasks, in a domain of cognitive functions (e.g. working memory); (3) variations in performance on the specific task at hand (e.g. “digit span” - listen to and repeat this sequence of numbers).

On its own, *g* does not address the boundary problems for the less general levels of this taxonomy: levels 2 and 3. As Deary notes concerning level 2, researchers disagree about the “nature of the domains—they can vary in number, name and content between samples depending on the battery applied—and there have long been worries about whether the nature of *g* might vary between cognitive batteries” [15]. These two less general levels are especially relevant for comparison with ML benchmarks: as we will argue in Section 4.1, task selection for ML benchmarks is value-laden in similar ways.

3.1.2 Boundaries and practical consequences influence each other

Technical considerations alone cannot resolve questions about the *significance* of the research. Judgments of significance in turn influence how we choose to define intelligence. The first IQ test, developed by Binet and Simon in 1905, was intended to help institutions identify students with learning difficulties, for the purpose of separating them from students of “normal” intelligence [9, 16, 50]. Other direct practical applications include labor and healthcare [16]. These applications help determine how the boundaries of the concept of intelligence are operationalized: what counts as intelligent behavior, and what does not. If we decide that, for example, intelligence should not be used to indicate health, that influences how we design tests to measure it. Thus, the values we attach to potential applications of a scientific concept influence how we define the boundaries of that concept.

Just as practical consequences influence boundaries, so do boundaries have practical consequences. The results of intelligence *research*, especially attempts at *explaining the causes* of differences in

⁸Current research attributes as much as nearly half of the variance in IQ scores to *g* [15].

test results, have a long history of practical uses, especially in justifying social hierarchies and structures of oppression [61, 10, 42, 15, 3].⁹ These practical consequences mean that choices made about the boundaries and significance of research have ethical ramifications. Choices made about the boundaries of intelligence research concepts can thus be value-laden in how they make certain ethical outcomes more likely than others. Similarly, in Section 4.1 we argue that choices about what tasks in ML are “significant” and what “domains” ML benchmarks should cover have ethical ramifications.

Just as importantly, as we see next (Section 3.2), intelligence researchers themselves *do and should* explicitly draw on ethical values in examining the validity of constructs in their research. This makes rejecting the place of ethical values in dealing with questions of boundaries especially implausible.

3.2 Validity and Values

Another important area where values play a role in intelligence research is in questions of validity [3]. How can we tell whether a given test, or a given statistical construct (e.g. *g*), measures something real, meaningful, or useful? Internal validity, which deals with the internal self-consistency of a concept, maps on to the idea of cross-validation accuracy in ML—the idea that a valid measurement of a construct should provide similar outcomes when applied to similar distributions of data. However, social scientists also define different types of *external* validity—types of validity that relate to how far a construct relates to phenomena outside its internal definition [59]. Here are a few examples.¹⁰

- (*Content validity*) How theoretically coherent is the construct being measured? Does operationalization of the construct agree with that theoretical understanding?
- (*Convergent validity*) Does the proposed measurement of a construct agree with other accepted measurements of the construct?
- (*Hypothesis validity*) How far are the measurements of the construct able to support substantively interesting hypotheses about the construct?
- (*Predictive validity*) Does the measure agree with properties that are coarsely related to our construct? For example, if we’re measuring IQ, does that correlate in expected ways with other properties and outcomes associated with intelligence?

These types of validity and others have been considered when evaluating human intelligence measures [76].

Even if we restrict our attention just to predictive validity, we still face choices of which types of predictive validity to favor. In Jensen [35]’s terms, IQ tests have validity if they “improve prediction of the quality of a person’s performance in a larger, more important sphere of activity”, if they can predict outcomes “that people deem important” [3]. This can be understood as a subset of external validity—where the concern is whether there are reliable connections between IQ measurements and phenomena outside of the scope of IQ tests. This common strategy for establishing that mental tests reliably relate to outcomes that are independently meaningful and useful turns on ethical values: health, educational, and occupational outcomes matter because people value them¹¹. Determining which types of validity are required to accept a benchmark like IQ, though, is a value-laden choice. It depends on how much we care about the external concepts that IQ is purportedly related to.

We believe that taking into account ethical considerations in choosing standards of evidence is something that intelligence researchers *ought to do*, rather than a mere mistake. We see this as an instance of what is sometimes called the *argument from inductive risk*: when the social costs of errors are

⁹Cave [10] helpfully explores the lessons of this family of issues around the practical uses of IQ—including in legitimating forms of domination—for the AI space in general. Anderson [3] points out an important historical twist to the role of intelligence research in justifying social hierarchies. Following Gould [30], Anderson argues that a case can be made that the historical debate between Spearman’s *g* and its detractors (including Thurstone) was in part motivated by disagreement about whether there is a single social hierarchy of abilities.

¹⁰See [33] for a more comprehensive overview. See [76] for a helpful example of how recent IQ tests tackle issues of validity.

¹¹In discussing the validity of IQ, intelligence researchers frequently emphasize that: (a) IQ predicts health and longevity outcomes [28, 36, 29, 71, 4, 72, 72]; (b) IQ predicts *education achievement* [17, 60]; (c) IQ predicts *occupational achievement*[18, 68, 60].

especially high, ethical, social, and political values *should* influence scientific standards of evidence [23, 67].

3.3 Path dependence

We want to consider one last respect in which intelligence research is value-laden: through *path dependence*—mechanisms that “lock in” historical antecedents and raise the switching cost of their alternatives [44, 57, 27]. To paraphrase Anderson, IQ has different ways of becoming a “self-fulfilling prophecy” [3].

One aspect of IQ and intelligence research that exhibits path dependence has to do with the real-world outcomes of intelligence research and IQ tests. Intelligence research and IQ tests influence behavior. An especially important area of influence concerns historically marginalized groups, including racial and gender disparities in IQ scores.¹² Take the case of the so-called “black-white IQ gap” in the US [34, 3]. Antiracist researchers have examined the role of mechanisms like *teacher expectation* and *stereotype threat* in reinforcing and entrenching historical disparities in cross-group test scores. For instance, false teacher beliefs about the lower potential of black students might lead students to “discouragement and disengagement from academic achievement” [26, 3]. Crucially, this sort of positive feedback loop makes IQ seem like more of a “real thing” that can be measured. If IQ tests are used to sort students into groups that receive different resources, those resources can influence how students do on later cognitive tests, thus making it seem like IQ caused those real-world outcomes, even if it was the differential allocation of resources or stereotypes that was really responsible.

Intelligence research doesn’t necessarily have to reinforce racism: antiracist researchers have shown how we can use research to discover mechanisms that reinforce path dependence, and to examine how such mechanisms can be countered and resisted. Whether or not to focus intelligence research in directions that reinforce or help undermine inequities is a value-laden choice.

4 Values in ML Evaluation

Having established that the debates over the right measures for human intelligence were unavoidably value-laden, we now seek to extend the same point to measures of artificial intelligence. Just as the concept of (human) intelligence is a thick evaluative concept (see Section 3), the concept of artificial intelligence (in the gradable sense of how intelligent a computer system is) is also a thick evaluative concept.¹³

The most direct parallels between human intelligence research and ML benchmarks lie in their respective attempts at setting standards for describing, evaluating, and comparing how different models or persons perform on tasks or collections of tasks.¹⁴ ML benchmarks need to enable comparisons of model performance in ways that involve both description and evaluation:

1. Through conveying commensurable empirical facts about the performance of different systems in a way that enables comparisons of model performance; and
2. Through conveying commensurable content about how we evaluate the systems (approval or disapproval, praise or blame, success or failure, etc.), in a way that enables ranking of model performance.

Evaluations of ML systems fulfill criterion 1: when researchers evaluate AI systems, they imagine themselves to be discovering empirical facts about how systems of a certain type behave. They also satisfy criterion 2: when we evaluate an ML system as being better or worse on a certain benchmark, we are evaluating it as being more or less successful. Given that ML evaluation has both descriptive and evaluative components, we can expect similar value-laden issues around ML evaluation to

¹²Here, we focus on race. For an overview of research on gender disparities in IQ scores, see Deary [14].

¹³Since AI systems often aim to mimic human cognitive abilities, this should be no surprise.

¹⁴An important difference between ML benchmarks and human intelligence difference is the latter’s interest in the environmental and hereditary causes of differences in intelligence. See 3.1 for our brief overview of the human intelligence case, and Deary [15] for a more detailed overview of different causal questions in intelligence research.

emerge, parallel to the issues we discussed with respect to measuring human intelligence. We now describe a few ways that ethical values and risks can enter ML benchmarks.

4.1 Task choices, task scopes, and task weightings are value-laden

The groundwork for making this claim has been laid by Raji et al. [58], who argue that the idea of a “universal” benchmark that is appropriate for evaluating AI across all contexts is unattainable and unhelpful. They advocate instead for benchmarks that are defined for particular tasks and contexts. If we accept their point that there is no such thing as a “general” task on which all AI systems should be evaluated, then the question arises of which tasks we should use to evaluate AI. The space of possible tasks is at the very least extremely large and possibly infinite, so it’s not feasible to evaluate AI on *all* tasks. We argue that values and interests unavoidably come into play when we make this choice of tasks, or when we choose to take performance on some tasks as more important than performance on other tasks. Ensmenger [25] makes similar points about the historical prevalence of chess as a game to evaluate AI against, arguing that it has shaped how AI is developed in value-laden ways.

These points parallel Anderson’s argument that even if we grant that there are value-neutral scientific facts, the choice of which of these facts are significant enough to seek out is value-laden [3]. Given limited scientific resources, we cannot possibly seek out all undiscovered facts out there. Choices about which truths and facts to pursue should be partly informed by their ethical, social and political implications.

Raji et al. [58]’s point is also analogous to debates in measuring human intelligence about whether there is a single general factor that explains variations in test scores, and if so, how it should be construed (see Section 3.1). Debates over which tasks to include in cognitive tests for human intelligence are parallel to debates over which tasks to include in ML benchmarks. In both cases the selection or weighting of tasks is made with reference to what we want our measurements to do outside of the research environment. What abilities do we want to incentivize AI systems to have? What abilities do we think are not worth incentivizing? These are ultimately questions that cannot be answered without referring to our ethical values around what type of AI we consider socially good and worth pursuing. As positive first steps, natural language processing (NLP) researchers have started commenting on the heavy emphasis on Indo-European languages in NLP datasets [32], and efforts to create benchmarks in previously neglected languages are under way [1].

4.2 Choices of validity in ML

In Section 3.2 we discussed how choices about which types of construct validity to use are value-laden. Here, we argue that similar concerns apply to ML benchmarks.

Typically, when they use cross-validation accuracy as the main metric of performance, ML researchers are measuring only internal validity, since the examples being validated on are randomly selected from the same dataset that the training data came from [43]. In contrast, the types of validity described in Section 3.2 are forms of external validity—relating the findings of a study to external phenomena outside of the study’s dataset. Focusing solely or mostly on internal validity via cross-validation is common when evaluating many types of ML models. In benchmarks that test only for internal validity, choosing not to consider external validity is a value-laden choice.

Some ML methods, like one-shot learning, are inherently evaluated in ways that go beyond internal validity, because they are evaluated on data distributions that are different from the training distribution. However, this is just *one* type of predictive validity, which is itself one type of external validity among many (see Section 3.2). Even when ML researchers do decide to measure external validity, there are further value-laden choices to be made about which types of external validity to measure. In Section 3.2 we discussed how external validity can be about the content of what’s measured (content validity), the convergence of the measurement with other ways of measuring the same construct (convergent validity), and whether the measurement correlates in the expected ways with other phenomena (predictive validity). Even *within* predictive validity, there are different targets that one can select as the phenomenon to find correlations with. Just as IQ researchers had to decide if IQ should predict health, education, or occupational achievements (among many other choices), ML researchers have to decide what types of external tasks or environments, they want their models to best succeed at. In other words, given that an ML researcher is looking for their model to do well in

environments or tasks outside the training data distribution, there are still further choices to be made about *which* environments or tasks these should be. Following Raji et al. [58]’s argument that there is no such thing as an “everything in the whole wide world” benchmark, it is not feasible for models to correlate equally well with *everything* in the external world. The choices made about which external phenomena we want models’ results to best correlate with inevitably depend on how we differentially value different external phenomena. For example, what novel language environments do we want models to be able to process inputs for? What counts as a “natural” photo or video for the purposes of a benchmark dataset (what external phenomena count as “natural”) [62]? How similar should machine vision be to human vision? What novel physical environments do we want a robot to be able to navigate? Liao et al. [43] emphasize the challenges in ML with establishing whether “progress on a benchmark transfers to other problems”—we argue here that the choice of *which* problems to measure these transfers on is inevitably value-laden, because it depends on what we want models to do in the real world, which in turn depends on our values.

Finally, failing to carefully consider external validity or selecting an inappropriate form of external validity can have direct ethical consequences. Many high-profile ML failures can be construed as failures of external validity [37, 20, 51]. Many critiques of ML attempts to “predict” properties like “criminality” or emotions can be interpreted as criticisms of the external validity of these constructs—for example, that criminality is a construct whose content cannot be captured in the way some ML practitioners attempt to (i.e. it lacks content validity) [12], or that emotions are defined in a psychologically implausible way by emotion detection algorithms (violating content validity and likely convergent validity as well) [66].

4.3 Path dependence

Just as measures of human intelligence can become self-fulfilling prophecies (Section 3.3), we argue that similar self-fulfilling dynamics can occur in AI evaluation. Intelligence tests for humans can be used to place certain humans in lower-resource environments that then cause them to do poorly on future intelligence tests. Similarly, ML benchmarks can discourage work on certain types of models, which in turn causes these types of models to do poorly (according to similar benchmarks) in the future. Current ML benchmarks are dominated by certain types of models—to wit, transformers and other types of deep neural networks [52, 54, 55, 53, 56]. As some have argued, the availability of certain types of data, the exigencies of current hardware, and the cultural prominence of certain types of problems among AI developers are at least partially responsible for this [31, 25, 7, 22]. Just as the practice of using *Drosophila* as a model organism in biology influenced the paths that 20th century biology took, Ensmenger [25] has argued that the preeminence of chess-playing as an ML problem has influenced how ML developed.

Path dependence can reinforce itself. Sara Hooker examines hardware-software interdependence as another source of self-reinforcing dynamics: some types of hardware, suitable to certain types of machine learning models or problems, have a healthy software ecosystem that can run on them. Other less dominant types of hardware also tend to have less developed software ecosystems associated with them. This puts the latter at an increasing disadvantage [31]. Hooker identifies multiple instances in computer science history of when “a research idea wins because it is suited to the available software and hardware and not because the idea is superior to alternative research directions.”

We can envision a similar feedback loop between benchmarks and models. As long as performance on currently popular benchmarks is what directs investment in modeling approaches that do well on those benchmarks, other modeling approaches will be increasingly at a disadvantage as they suffer from lack of computational resources, hardware incompatibilities, and lack of researcher interest [27]. For example, we do not know if Bayesian models could achieve similar performances on the same problems if we were to pour the same amount of resources into them that we do for neural networks: resources to improve their architectures, fine-tune them, and train them on large amounts of data. Compared to some other modeling approaches, neural networks also get comparatively more investment in the creation and maintenance of sophisticated, open source software packages like Tensorflow and Pytorch. Relying on benchmarks that largely favor neural networks could give this feedback loop an additional boost. The much-criticized practice of “SOTA chasing” (creating models with the main purpose of getting the best “state-of-the-art” performances on benchmarks) [11] can be construed as a form of path dependence: it encourages investment into incremental improvements on dominant architectures and modeling approaches, where in a different world, we

might instead have more investment in less popular model types that are harder to work with given the current software/hardware environment.

In a different world, one can also imagine having an evaluation environment that corrects for some of the disadvantages of non-dominant model types. For example, some NLP researchers have called for performance measures that adjust for the extent of hyperparameter tuning and computational budget [21]. The fact that doing so is not currently widespread practice is a value-laden choice that puts non-dominant model architectures at a disadvantage. However, even if we did correct for computational resources and fine-tuning, we cannot correct for the problem that far more researcher-hours are spent improving neural networks and exploring their possibilities, with comparatively much less effort being spent on other alternatives. As a corrective, research funds could be distributed with more attention to the need to not over-focus on already-popular approaches—but whether to do this or not is a value-laden choice that, given current uncertainties about which path is best, cannot be decided by purely technical considerations.

Once we acknowledge path dependence as a real phenomenon, and recognize that choosing certain benchmarks can make the playing field less level because of how it influences future investment in AI, it becomes crucial to recognize this as an area of value-laden decisions. We frame this as a choice between the following:

1. Validating model performance against current benchmarks, assuming current distributions of resources for the development of different types of models.
2. Validating model performance in a hypothetical domain where all model types have been tuned to the same extent, use similar computational resources, and have received similar investments in architectural improvements and software ecosystems.

We do not intend here to insist that one or the other choice is the *correct* one, but to point out that choosing one or the other is not merely a matter of fact or straightforwardly dictated by the “scientific method”. Ethical, social and political values ought to inform our choices here, because the homogenization of the model landscape has ethical ramifications and risks [27, 40, 13].

4.3.1 Path dependence in reifying social constructs

Another form of path dependency is when social constructs used as categories in ML classification benchmarks are themselves made “more real” or more persistent by the fact that they are being used in benchmarks. In Section 3.3, we explained how using IQ tests to differentially allocate resources to students can make IQ itself seem like more of a “real” measure. The differential allocation of resources based on IQ can lead to systematic differences in students’ later achievements in life, thus giving IQ an apparent external validity. Similarly, the increasing use of ML to, for example, classify images of humans into gender categories can reinforce the very same visual norms about gender categories that the ML models presume to be objective [63]. Using benchmarks that contain such categories encourages ML practitioners to develop models that do well on those benchmarks—which means the models have to do well at mimicking the categories presumed in the benchmark’s “ground truth” dataset. These models can then be used in the wider world to classify real humans into gender categories, which means that real humans may be motivated to conform to the visual norms of those categories in order to be classified “correctly” by the models. When real humans start modifying their behavior in this way, the norms of gendered visual presentations are strengthened, meaning that apparent gender categories in real-world images that future models will be trained on are also accentuated.

Thus, ML evaluation runs the risk of reifying various social categories or norms that would otherwise be less rigid. Deciding whether or not to include those types of categorizations in our benchmarks is therefore a value-laden choice that depends on whether we think those are good categories to reify. Previous work has criticized current practices around gender categories in ML [63]. Many other social categories like race are also value-laden in the same way. Although the creators of benchmarks like ImageNet might think that they are simply doing a value-neutral catalog of images of “the world” [58], their choices of image labels as deeply value-laden, as others have argued [19, 64].

5 Practical recommendations

Being reflective, explicit, and public about the social, political, and ethical values behind ML research is vital to the pursuit of responsible ML. It is also paramount for any attempt at developing adequate ethics guidelines and ethics review for a premier ML research venue like NeurIPS.

We have argued that in at least three ways, scientific and technical decisions about ML benchmarks are at least implicitly dependent on ethical values. From the perspective of identifying and mitigating the ethical risks of ML benchmark research, the areas of task selection, choice of validity standards, and path dependence each need close and thorough scrutiny.

A first recommendation is for people working on ethics guidelines and ethics reviews for ML benchmarks in communities like NeurIPS’s Datasets and Benchmarks track. Given the unique role of ML benchmarks in enabling the evaluation and comparison of ML models, examination of the ethical risks of benchmarks should not be narrowly limited to considerations about their potential uses, or about the treatment of human subjects. Here, we’ve identified sources of ethical risk that come with seemingly technical aspects of benchmarks. Are there other key decision points about the technical features of benchmarks that also pose significant ethical risk? What actionable guidelines might enable researchers and reviewers to more reflectively and explicitly consider these ethical risks? We hope that the three areas we have identified are a helpful starting point in answering these questions.

A second recommendation is aimed at researchers. Given the unique role of benchmarks in ML research, researchers should avoid interpreting performance on benchmarks as a value-neutral indicator of capabilities. Moreover, we recommend that researchers developing benchmarks explicitly and reflectively discuss the potential ethical risk that come with areas like task selection, standards of validity, and path dependence, either through part of benchmark documentation, as part of impact statements, or as part of the papers themselves. This will in turn help other practitioners choose which benchmarks to use when evaluating their models. For example, if a particular form of external validity (e.g. consequential validity) is particularly important to a real-world application, data scientists who are evaluating models can more easily choose benchmarks that emphasize consequential validity, while putting less weight on benchmarks that don’t. This will better align model evaluation with the values the practitioner wants.

Finally, we want to acknowledge an important practical challenge in mitigating the ethical risks that come with the areas we have identified. Path dependence is an example of a *structural problem* that is unlikely to be resolved through isolated individual action. Mitigating the ethical risks that come with path dependence calls for social and collective change in areas like how research is funded and incentivized. For the ethical risks of benchmarks that stem from structural problems like path dependence, future work should investigate strategies and opportunities for levelling the playing field for diverse approaches to ML, in order to minimize path dependence.

6 Conclusion

What does *good science* look like in the context of ML benchmarks? We’ve argued that benchmarks are and should be influenced by ethical, social, and political values. The concept of intelligence is value-laden: we define what intelligence is according to how we want machines and people to perform. Moreover, decisions about benchmarks for ML and human intelligence are intimately tied to their uses. The practical implications of ML benchmarks and IQ mean that the social costs of errors can be very high.

The argument from inductive risk (see Section 3.2) applies to each of the dimensions of the problem we considered for both ML benchmark and human intelligence research. What tasks we want machines (and humans) to perform well, what standards of validity we want to bring to research on either, and decisions about where to curb or reinforce path dependence *should be* deeply informed by what purposes we want them to serve in human society.

We believe that for scientific research to be valuable, it must produce reliable empirical knowledge [24]. This requires prioritizing considerations about what is more or less conducive to the truth—about *epistemic values*—throughout. Key background decisions (about task selection, standards of validity, and whether to reinforce or curb path dependence) must be informed by both epistemic considerations, and by reflective and explicit articulation of the risks of research, and of the ethical,

social, and political values we want to prioritize in mitigating those risks. Yet once we have answered these background questions, we should use only epistemic values to evaluate research according to our chosen standards.¹⁵ Paraphrasing Douglas [24], intelligence is built on a social, political, and moral terrain. Good science for ML benchmarks and human intelligence research requires reflective and explicit articulation of the values and risks embedded in their scientific and technical core.¹⁶

References

- [1] Jade Abbott and Laura Martinus. 2019. Benchmarking Neural Machine Translation for Southern African Languages. In *Proceedings of the 2019 Workshop on Widening NLP*. Association for Computational Linguistics, Florence, Italy, 98–101. <https://aclanthology.org/W19-3632>
- [2] Anna Alexandrova and Mark Fabian. 2022. Democratising Measurement: or Why Thick Concepts Call for Coproduction. *European Journal for Philosophy of Science* 12, 1 (March 2022), 7. <https://doi.org/10.1007/s13194-021-00437-7>
- [3] Elizabeth Anderson. 2002. Situated Knowledge and the Interplay of Value Judgments and Evidence in Scientific Inquiry. In *In the Scope of Logic, Methodology and Philosophy of Science*, Peter Gärdenfors, Jan Woleński, and Katarzyna Kijania-Placek (Eds.). Springer Netherlands, Dordrecht, 497–517. https://doi.org/10.1007/978-94-017-0475-5_8
- [4] Rosalind Arden, Michelle Luciano, Ian J Deary, Chandra A Reynolds, Nancy L Pedersen, Brenda L Plassman, Matt McGue, Kaare Christensen, and Peter M Visscher. 2016. The association between intelligence and lifespan is mostly genetic. *International Journal of Epidemiology* 45, 1 (Feb. 2016), 178–185. <https://doi.org/10.1093/ije/dyv112>
- [5] Samy Bengio, Alina Beygelzimer, Kate Crawford, Jeanne Fromer, Iason Gabriel, Amanda Levendowski, Deborah Raji, and Marc’Aurelio Ranzato. 2022. Provisional Draft of the NeurIPS Code of Ethics. <https://openreview.net/forum?id=zVoy8kAFKPr>
- [6] Abeba Birhane, Pratyusha Kalluri, Dallas Card, William Agnew, Ravit Dotan, and Michelle Bao. 2022. The Values Encoded in Machine Learning Research. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Seoul Republic of Korea, 173–184. <https://doi.org/10.1145/3531146.3533083>
- [7] Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. 2020. Experience Grounds Language. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online, 8718–8735. <https://doi.org/10.18653/v1/2020.emnlp-main.703>
- [8] Rodrigo Canaan, Christoph Salge, Julian Togelius, and Andy Nealen. 2019. Leveling the playing field: Fairness in AI versus human game benchmarks. In *Proceedings of the 14th International Conference on the Foundations of Digital Games*. 1–8.
- [9] John B. Carroll. 1993. *Human cognitive abilities: a survey of factor-analytic studies*. Cambridge University Press, Cambridge ; New York.
- [10] Stephen Cave. 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*. ACM, New York NY USA, 29–35. <https://doi.org/10.1145/3375627.3375813>
- [11] Kenneth Ward Church and Valia Kordoni. 2022. Emerging trends: SOTA-chasing. *Natural Language Engineering* 28, 2 (2022), 249–269.
- [12] Amanda Coston, Anna Kawakami, Haiyi Zhu, Ken Holstein, and Hoda Heidari. 2022. A Validity Perspective on Evaluating the Justified Use of Data-driven Decision-making Algorithms. <https://doi.org/10.48550/ARXIV.2206.14983>

¹⁵This is similar to Anderson’s take on impartiality Anderson [3].

¹⁶For more on the conception of good science we recommend for benchmarks, see A.1.

- [13] Kathleen Creel and Deborah Hellman. 2022. The Algorithmic Leviathan: Arbitrariness, Fairness, and Opportunity in Algorithmic Decision-Making Systems. *Canadian Journal of Philosophy* (2022), 1–18.
- [14] Ian J. Deary. 2012. 125 Years of Intelligence in *The American Journal of Psychology*. *The American Journal of Psychology* 125, 2 (July 2012), 145–154. <https://doi.org/10.5406/amerjpsyc.125.2.0145>
- [15] Ian J. Deary. 2012. Intelligence. *Annual Review of Psychology* 63, 1 (Jan. 2012), 453–482. <https://doi.org/10.1146/annurev-psych-120710-100353>
- [16] Ian J. Deary. 2020. *Intelligence: A Very Short Introduction*. Oxford University Press. <https://www.veryshortintroductions.com/view/10.1093/actrade/9780198796206.001.0001/actrade-9> OCLC: 1181910522.
- [17] Ian J. Deary, Steve Strand, Pauline Smith, and Cres Fernandes. 2007. Intelligence and educational achievement. *Intelligence* 35, 1 (Jan. 2007), 13–21. <https://doi.org/10.1016/j.intell.2006.02.001>
- [18] Ian J. Deary, Michelle D. Taylor, Carole L. Hart, Valerie Wilson, George Davey Smith, David Blane, and John M. Starr. 2005. Intergenerational social mobility and mid-life status attainment: Influences of childhood intelligence, childhood social factors, and education. *Intelligence* 33, 5 (Sept. 2005), 455–472. <https://doi.org/10.1016/j.intell.2005.06.003>
- [19] Emily Denton, Alex Hanna, Razvan Amironesei, Andrew Smart, and Hilary Nicole. 2021. On the genealogy of machine learning datasets: A critical history of ImageNet. *Big Data & Society* 8, 2 (2021).
- [20] Ben Dickson. 2019. Human Help Wanted: Why AI Is Terrible at Content Moderation. <https://www.pcmag.com/opinions/human-help-wanted-why-ai-is-terrible-at-content-moderation> Accessed Jan 18, 2022.
- [21] Jesse Dodge, Suchin Gururangan, Dallas Card, Roy Schwartz, and Noah A. Smith. 2019. Show Your Work: Improved Reporting of Experimental Results. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Hong Kong, China, 2185–2194. <https://doi.org/10.18653/v1/D19-1224>
- [22] Ravit Dotan and Smitha Milli. 2020. Value-laden disciplinary shifts in machine learning. *FAT* '20: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Jan. 2020). <https://doi.org/10.1145/3351095.3373157>
- [23] Heather Douglas. 2000. Inductive Risk and Values in Science. *Philosophy of Science* 67, 4 (Dec. 2000), 559–579. <https://doi.org/10.1086/392855>
- [24] Heather Douglas. 2014. The Moral Terrain of Science. *Erkenntnis* 79, S5 (June 2014), 961–979. <https://doi.org/10.1007/s10670-013-9538-0>
- [25] Nathan Ensmenger. 2012. Is chess the drosophila of artificial intelligence? A social history of an algorithm. *Social studies of science* 42, 1 (2012), 5–30.
- [26] Ronald F Ferguson. 2003. Teachers’ perceptions and expectations and the Black-White test score gap. *Urban education* 38, 4 (2003), 460–507.
- [27] Nic Fishman and Leif Hancox-Li. 2022. Should attention be all we need? The epistemic and ethical implications of unification in machine learning. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. <https://doi.org/10.1145/3531146.3533206>
- [28] Linda S. Gottfredson. 1997. Why g matters: The complexity of everyday life. *Intelligence* 24, 1 (Jan. 1997), 79–132. [https://doi.org/10.1016/S0160-2896\(97\)90014-3](https://doi.org/10.1016/S0160-2896(97)90014-3)

- [29] Linda S. Gottfredson and Ian J. Deary. 2004. Intelligence Predicts Health and Longevity, but Why? *Current Directions in Psychological Science* 13, 1 (Feb. 2004), 1–4. <https://doi.org/10.1111/j.0963-7214.2004.01301001.x>
- [30] Stephen Jay Gould. 1981. *The mismeasure of man* (1st ed.). Norton, New York.
- [31] Sara Hooker. 2021. The hardware lottery. *Commun. ACM* 64, 12 (2021), 58–65.
- [32] Dirk Hovy and Shannon L. Spruit. 2016. The Social Impact of Natural Language Processing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, Berlin, Germany, 591–598. <https://doi.org/10.18653/v1/P16-2096>
- [33] Abigail Z. Jacobs and Hanna Wallach. 2021. Measurement and Fairness. *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (March 2021), 375–385. <https://doi.org/10.1145/3442188.3445901>
- [34] Christopher Jencks and Meredith Phillips (Eds.). 1998. *The Black–White test score gap*. Brookings Institution Press, Washington, DC, US. Pages: x, 523.
- [35] Arthur Robert Jensen. 1981. *Straight talk about mental tests*. Free Press, New York.
- [36] Arthur Robert Jensen. 1998. *The g factor: the science of mental ability*. Praeger, Westport, Conn.
- [37] Christopher J Kelly, Alan Karthikesalingam, Mustafa Suleyman, Greg Corrado, and Dominic King. 2019. Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine* 17, 1 (2019), 1–9.
- [38] Simon Kirchin. 2013. Introduction: Thick and Thin Concepts. In *Thick concepts* (1st ed.), Simon Kirchin (Ed.). Oxford University Press, Oxford.
- [39] Simon Kirchin. 2017. *Thick evaluation* (1st ed.). Oxford University Press, Oxford.
- [40] Jon Kleinberg and Manish Raghavan. 2021. Algorithmic monoculture and social welfare. *Proceedings of the National Academy of Sciences* 118, 22 (2021). <https://doi.org/10.1073/pnas.2018340118>
- [41] Thomas S. Kuhn. 1977. Objectivity, Value Judgment, and Theory Choice. In *The essential tension: selected studies in scientific tradition and change*. Univ. of Chicago Press, Chicago, Ill., 320–339.
- [42] Philippa Levine. 2017. *Eugenics: a very short introduction*. Oxford University Press, New York, NY.
- [43] Thomas Liao, Rohan Taori, Deborah Raji, and Ludwig Schmidt. 2021. Are We Learning Yet? A Meta Review of Evaluation Failures Across Machine Learning. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, J. Vanschoren and S. Yeung (Eds.), Vol. 1. <https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/file/757b505cfd34c64c85ca5b569>
- [44] Stan J Liebowitz and Stephen E Margolis. 1995. Path dependence, lock-in, and history. *Journal of Law, Economics, & Organization* 11 (1995), 205–226. Issue 1.
- [45] Yusen Liu, Fangyuan He, Haodi Zhang, Guozheng Rao, Zhiyong Feng, and Yi Zhou. 2019. How Well Do Machines Perform on IQ tests: a Comparison Study on a Large-Scale Dataset. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, Macao, China, 6110–6116. <https://doi.org/10.24963/ijcai.2019/846>
- [46] Helen E. Longino. 1996. Cognitive and Non-Cognitive Values in Science: Rethinking the Dichotomy. In *Feminism, Science, and the Philosophy of Science*, Lynn Hankinson Nelson and Jack Nelson (Eds.). Kluwer Academic Publishers, 39–58.

- [47] Sasha Luccioni, William Isaac, Cherie Poland, Deborah Raji, Samy Bengio, Kate Crawford, Jeanne Fromer, Iason Gabriel, Amanda Levendowski, and Marc’Aurelio Ranzato. 2021. Ethical Review Guidelines. <https://neurips.cc/public/EthicsGuidelines> Accessed Aug 28, 2022.
- [48] Gary Marcus and Ernest Davis. 2019. *Rebooting AI: Building artificial intelligence we can trust*. Pantheon, New York.
- [49] Ernan McMullin. 1982. Values in Science. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1982, 4 (1982), 3–28. <https://doi.org/10.1111/j.1467-9744.2012.01298.x> Publisher: Philosophy of Science Association.
- [50] Kevin J. Mitchell. 2018. *Innate: how the wiring of our brains shapes who we are*. Princeton University Press, Princeton.
- [51] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447–453. <https://doi.org/10.1126/science.aax2342> arXiv:<https://www.science.org/doi/pdf/10.1126/science.aax2342>
- [52] Papers With Code. 2022. Image Classification on ImageNet. <https://paperswithcode.com/sota/image-classification-on-imagenet> Accessed Apr 22, 2022.
- [53] Papers With Code. 2022. Machine Translation on WMT2014 English-German. <https://paperswithcode.com/sota/machine-translation-on-wmt2014-english-german> Accessed Apr 22, 2022.
- [54] Papers With Code. 2022. Object Detection on COCO test-dev. <https://paperswithcode.com/sota/object-detection-on-coco> Accessed Apr 22, 2022.
- [55] Papers With Code. 2022. Semantic Segmentation on ADE20K. <https://paperswithcode.com/sota/semantic-segmentation-on-ade20k> Accessed Apr 22, 2022.
- [56] Papers With Code. 2022. Speech Recognition on LibriSpeech test-clean. <https://paperswithcode.com/sota/speech-recognition-on-librispeech-test-clean> Accessed Apr 22, 2022.
- [57] Mark S. Peacock. 2009. Path Dependence in the Production of Scientific Knowledge. *Social Epistemology* 23, 2 (April 2009), 105–124. <https://doi.org/10.1080/02691720902962813>
- [58] Inioluwa Deborah Raji, Emily M Bender, Amandalynne Paullada, Emily Denton, and Alex Hanna. 2021. AI and the everything in the whole wide world benchmark. *Conference on Neural Information Processing Systems (NeurIPS 2021) Track on Datasets and Benchmarks* (2021).
- [59] Julian Reiss. 2019. Against external validity. *Synthese* 196, 8 (Aug. 2019), 3103–3121. <https://doi.org/10.1007/s11229-018-1796-6>
- [60] Stuart Ritchie. 2015. *Intelligence: all that matters*. Hodder & Stoughton, London.
- [61] Angela Saini. 2019. *Superior: the return of race science*. Beacon Press, Boston.
- [62] Morgan Klaus Scheuerman, Alex Hanna, and Emily Denton. 2021. Do datasets have politics? Disciplinary values in computer vision dataset development. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–37.
- [63] Morgan Klaus Scheuerman, Jacob M. Paul, and Jed R. Brubaker. 2019. How Computers See Gender: An Evaluation of Gender Classification in Commercial Facial Analysis Services. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 144 (Nov. 2019), 33 pages. <https://doi.org/10.1145/3359246>

- [64] Morgan Klaus Scheuerman, Kandrea Wade, Caitlin Lustig, and Jed R. Brubaker. 2020. How We’ve Taught Algorithms to See Identity: Constructing Race and Gender in Image Databases for Facial Analysis. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1, Article 58 (May 2020), 35 pages. <https://doi.org/10.1145/3392866>
- [65] C. Spearman. 1904. "General Intelligence," Objectively Determined and Measured. *The American Journal of Psychology* 15, 2 (April 1904), 201. <https://doi.org/10.2307/1412107>
- [66] Luke Stark and Jesse Hoey. 2021. The ethics of emotion in artificial intelligence systems. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 782–793.
- [67] Daniel Steel. 2010. Epistemic Values and the Argument from Inductive Risk. *Philosophy of Science* 77, 1 (Jan. 2010), 14–34. <https://doi.org/10.1086/650206>
- [68] Tarmo Strenze. 2007. Intelligence and socioeconomic success: A meta-analytic review of longitudinal research. *Intelligence* 35, 5 (Sept. 2007), 401–426. <https://doi.org/10.1016/j.intell.2006.09.004>
- [69] Pekka Väyrynen. 2021. Thick Ethical Concepts. In *The Stanford Encyclopedia of Philosophy* (Spring 2021 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University.
- [70] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461* (2018).
- [71] Russell T. Warne. 2016. Five Reasons to Put the *g* Back Into *G* iftleness: An Argument for Applying the Cattell–Horn–Carroll Theory of Intelligence to Gifted Education Research and Practice. *Gifted Child Quarterly* 60, 1 (Jan. 2016), 3–15. <https://doi.org/10.1177/0016986215605360>
- [72] Russell T. Warne. 2021. Between-Group Mean Differences in Intelligence in the United States Are >0% Genetically Caused: Five Converging Lines of Evidence. *The American Journal of Psychology* 134, 4 (Dec. 2021), 480–501. <https://doi.org/10.5406/amerjpsyc.134.4.0479>
- [73] Russell T. Warne and Cassidy Burningham. 2019. Spearman’s *g* found in 31 non-Western nations: Strong evidence that *g* is a universal phenomenon. *Psychological Bulletin* 145, 3 (March 2019), 237–272. <https://doi.org/10.1037/bul0000184>
- [74] Max Weber. 1917. Der Sinn der “Wertfreiheit“ der soziologischen und ökonomischen Wissenschaften. *Logos. Internationale Zeitschrift für Philosophie der Kultur* VII (1917).
- [75] Max Weber. 2011. The Meaning of “Ethical Neutrality” in Sociology and Economics. In *Methodology of Social Sciences*. Routledge. <https://www.taylorfrancis.com/books/e/9781315124445> OCLC: 1004359569.
- [76] David Wechsler. 2008. *Wechsler adult intelligence scale: WAIS-IV ; technical and interpretive manual* (4th ed.). Pearson, San Antonio, Texas.
- [77] Langdon Winner. 1978. *Autonomous technology: technics-out-of-control as a theme in political thought*. MIT Press, Cambridge, Massachusetts.

A Appendix

A.1 Value neutrality and feminist philosophy of science

Our perspective in this paper is deeply informed by feminist philosophy of science scholarship about the value neutrality of science. In this appendix, we provide background on the conception of value neutrality implicit in this paper (A.1.1). We also argue that feminist critiques of value neutrality are a vital source of insights for what ML research, including for articulating the place of ethical, social, and political values in *good science* (A.1.2). We mean this as an invitation for other researchers to dig further into the many ways this deep body of scholarship can be helpful to the ML research context.

A.1.1 What is value-neutrality?

Discussions of the value neutrality of science have a long history. McMullin [49] interprets Weber as having argued—over a hundred years ago—that "the objectivity of science [...] requires public norms accessible to all, and interpreted by all in the same way" [74, 75]. On this early 20th century view, values that are open to choice and disagreement have no place in science.

By contrast, for multiple decades, philosophers have recognized a role for values that are open to dispute and choice in scientific reasoning and decisions [41, 49, 46, 22].¹⁷ A particularly important concern is the *underdetermination thesis*: scientific theory is underdetermined by empirical evidence [46, 49]. Empirical data cannot on its own settle the question of what specific theory is correct. A commonly accepted view is that values have a role to play in closing the gap between empirical data and theory choice.

This hasn't spelled the end of the value neutrality of science. We can hold on to the ideal of value-neutral science, some argue, if the only values we rely on in closing the gap between empirical data and theory choice are *epistemic values*: that is, considerations about what more or less promotes the attainment of truth [67, 46].¹⁸ Epistemically problematic practices like falsifying records or cherry-picking data undermine the core functions that make science worth pursuing, such as the aim of producing "reliable empirical knowledge" [24]. Acknowledging a role for epistemic values in scientific reasoning is an invitation to clarify specifically what we consider to promote the attainment of truth or of reliable empirical knowledge. Kuhn's influential account highlights 5 key epistemic values: accuracy, simplicity, internal and external consistency, breadth of scope, and fruitfulness [41, 46].

Of note, this philosophy of science sense of *value neutrality* is different from the sense of value-neutrality centered in Birhane et al. [6]'s helpful investigation of the values of ML research. Their account instead center's Winner [77]'s conception of the value-neutrality: the view that technology is value neutral if it can be put to both beneficial and harmful uses, and if whether it is harmful or not harmful depends on what uses and purposes we choose to put it to. This is the kind of "neutrality" that household tools have. A hammer can be used to attach a work of art to a wall; it can also be used as a weapon.

The philosophy of science sense of value neutrality is instead about what kinds of values are admissible in scientific reasoning. Contemporary versions of the value neutrality of science thesis allow *epistemic values* to supplement empirical evidence in closing the gap between evidence and theory. What the value-neutrality thesis rejects is any place for ethical, social, or political values (non-epistemic values more broadly) in scientific reasoning.

A.1.2 The relevance of feminist philosophy of science for ML research

There are two aspects of feminist philosophers of science's rejection of value neutrality that we find especially helpful for the ML space.

The first concerns the status of scientific research that purports or appears to be *value-neutral*—that does not make explicit and reflective appeal to ethical, social, and political values. In this paper, we have argued that IQ and ML benchmark research, even in cases where it purports or appears to be value-neutral, nonetheless embodies ethical, social and/or political values. We see this as deeply aligned with a growing body of scholarship on the values embodied in ML research [22, 27, 6]. In doing so, we hope to spark interest in examining the many other ways in which insights and models from the feminist philosophy of science tradition can help reconsider the many roles of unacknowledged and implicit ethical, social, and political values in ML research.

¹⁷For a helpful account connecting philosophy of science to values in ML research, see Dotan and Milli [22]. On the philosophy of science front, our accounts center different bodies of work: "Kuhn's paradigms and Lakatos's research programmes" in their case; feminist scholarship on the value neutrality of science, epistemic/non-epistemic values, thick evaluative concepts in social science, and the argument from inductive risk in ours.

¹⁸In some 90's discussions, epistemic values are instead called "cognitive values". Recent discussions favor the term epistemic values to avoid presupposing (or seeming to presuppose) non-cognitivism about ethical, social, and political values.

The second concerns the status of research that is explicitly and reflectively informed by ethical, social, and political values, like the insights of anti-racist and feminist IQ research that Elizabeth Anderson examines (discussed above in Section 3). Inspired by feminist standpoint epistemology’s emphasis on the epistemic privilege of the standpoints of oppressed and historically marginalized communities, feminist philosophers of science have been investigating the multiple ways in which explicit and reflective reliance on ethical, social, and political values can make for *better science* [3]. As discussed in Section 6, this perspective is compatible with a strong endorsement of the crucial role of empirical evidence in science, and of the distinct value of science in producing reliable empirical knowledge, and of the impartiality of science.

A lot of research on ethical and just ML deeply aligns with this perspective. We believe that the feminist philosophy of science tradition is an especially valuable source of models for the multiple ways in which *good science* is and should be informed by ethical, social, and political values.

We want to conclude by highlighting an ideal for good science that we especially hope to see explored by the ML benchmark community. Once we accept that benchmarking intelligence is value-laden—that ethical, social and political values should inform how we select standards for describing, evaluating, and comparing human or machine intelligence—we can begin to consider how to define benchmarks in order to promote the values that we want. In the case of measuring human intelligence, Anderson [3] suggests explicitly taking up the *epistemological standpoint of justice*. This is an injunction to: (a) focus the efforts of research on figuring out what it would take for members of all groups to fully develop their potential—as opposed to today’s society, where members of vulnerable and historically marginalized communities are systematically deprived of opportunities to do so; and (b) to thoroughly investigate what obstacles (e.g. path dependency) get in the way. It’s an epistemological standpoint because our response to the injunction should turn on “discoverable empirical facts” and reliable empirical knowledge, rather than merely on non-epistemic values. In the case of human intelligence, examples include studying mechanisms like teacher expectation and stereotype threat (Section 3.3), and experimenting with social reforms [3].

What would it take to adopt the epistemological standpoint of justice for ML? For example, what measures of machine intelligence would enable different groups of humans to thrive in a society where they interact regularly with ML systems? What measures of machine intelligence might reinforce harms to vulnerable and historically marginalized communities (e.g. through path dependence and the reification of social constructs)? Misguided belief in “neutral” approaches to evaluating intelligence—as though we can simply measure “everything in the whole wide world”—makes for less useful and less significant science.