

Can We Trust AI Benchmarks?

An Interdisciplinary Review of Current Issues in AI Evaluation

MARIA ERIKSSON*, European Commission, Joint Research Centre (JRC), Seville, Spain

ERASMO PURIFICATO*, European Commission, Joint Research Centre (JRC), Ispra, Italy

ARMAN NOROOZIAN*, European Commission, Joint Research Centre (JRC), Brussels, Belgium

JOÃO VINAGRE*, European Commission, Joint Research Centre (JRC), Seville, Spain

GUILLAUME CHASLOT*, European Commission, Joint Research Centre (JRC), Brussels, Belgium

EMILIA GOMEZ*, European Commission, Joint Research Centre (JRC), Seville, Spain

DAVID FERNANDEZ-LLORCA*, European Commission, Joint Research Centre (JRC), Seville, Spain

Quantitative Artificial Intelligence (AI) Benchmarks have emerged as fundamental tools for evaluating the performance, capability, and safety of AI models and systems. Currently, they shape the direction of AI development and are playing an increasingly prominent role in regulatory frameworks. As their influence grows, however, so too does concerns about *how* and with what effects they evaluate highly sensitive topics such as capabilities, including high-impact capabilities, safety and systemic risks. This paper presents an interdisciplinary meta-review of about 100 studies that discuss shortcomings in quantitative benchmarking practices, published in the last 10 years. **It brings together many fine-grained issues in the design and** application of benchmarks (such as biases in dataset creation, inadequate documentation, data contamination, and failures to distinguish signal from noise) with broader sociotechnical issues (such as an over-focus on evaluating text-based AI models according to one-time testing logic that fails to account for how AI models are increasingly multimodal and interact with humans and other technical systems). **Our review also highlights a series of systemic flaws in current** benchmarking practices, such as **misaligned incentives, construct validity issues, unknown unknowns**, and problems with the **gaming of benchmark results**. Furthermore, it underscores how benchmark practices are fundamentally shaped by cultural, commercial and competitive dynamics that often prioritise state-of-the-art performance at the expense of broader societal concerns. By providing an overview of risks associated with existing benchmarking procedures, we problematise **disproportionate trust placed in benchmarks** and contribute to ongoing efforts to improve the accountability and relevance of quantitative AI benchmarks within the complexities of real-world scenarios.

CCS Concepts: • **General and reference** → **Evaluation; Metrics; Measurement; Validation; Reliability**.

Additional Key Words and Phrases: AI benchmarks, Benchmark critique, AI evaluation, Safety evaluation, AI Regulation

***Disclaimer:** The views expressed in this paper are purely those of the authors and may not, under any circumstances, be regarded as an official position of the European Commission.

Authors' Contact Information: [Maria Eriksson](mailto:maria.eriksson@ec.europa.eu), maria.eriksson@ec.europa.eu, European Commission, Joint Research Centre (JRC), Seville, Spain; [Erasmus Purificato](mailto:erasmo.purificato@ec.europa.eu), erasmo.purificato@ec.europa.eu, European Commission, Joint Research Centre (JRC), Ispra, Italy; [Arman Noroozian](mailto:arman.noroozian@ec.europa.eu), arman.noroozian@ec.europa.eu, European Commission, Joint Research Centre (JRC), Brussels, Belgium; [João Vinagre](mailto:joao.vinagre@ec.europa.eu), joao.vinagre@ec.europa.eu, European Commission, Joint Research Centre (JRC), Seville, Spain; [Guillaume Chaslot](mailto:guillaume.chaslot@ec.europa.eu), guillaume.chaslot@ec.europa.eu, European Commission, Joint Research Centre (JRC), Brussels, Belgium; [Emilia Gomez](mailto:emilia.gomez-gutierrez@ec.europa.eu), emilia.gomez-gutierrez@ec.europa.eu, European Commission, Joint Research Centre (JRC), Seville, Spain; [David Fernandez-Llorca](mailto:david.fernandez-llorca@ec.europa.eu), david.fernandez-llorca@ec.europa.eu, European Commission, Joint Research Centre (JRC), Seville, Spain.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Manuscript submitted to ACM

Manuscript submitted to ACM

1

ACM Reference Format:

Maria Eriksson, Erasmo Purificato, Arman Noroozian, João Vinagre, Guillaume Chaslot, Emilia Gomez, and David Fernandez-Llorca. 2025. Can We Trust AI Benchmarks? An Interdisciplinary Review of Current Issues in AI Evaluation. 1, 1 (February 2025), 20 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Quantitative artificial intelligence (AI) benchmarks (i.e., combinations of test datasets and performance metrics that are taken to represent general or specific tasks and used to compare AI model capabilities and/or risks [87]) play a central role in the release and marketing of newly developed AI tools. Together with qualitative evaluation methods (such as red teaming and peer confrontations [20]), quantitative benchmarks - hereafter referred to as *AI benchmarks*, or simply *benchmarks* - are generally seen as providing crucial feedback signals on the performance and capabilities of AI. **Indeed, they have become so critical to AI development that businesses go to great lengths to achieve good benchmarking** scores, with market players like OpenAI being estimated to have spent hundreds of thousands of dollars on compute to obtain a high score at the ARC-AGI benchmark [84].

Increasingly, AI benchmarks are also used in regulatory contexts, where the goal is to assess potential societal harms posed by AI models and systems. The most notable case is the EU AI Act [30], which incorporates benchmarks in several key provisions. For instance, for high-risk AI systems, benchmarks are expected to play a significant role in complying with requirements on accuracy, robustness, and cybersecurity (Art. 15(2)), and their development will be facilitated through AI regulatory sandboxes (Art. 58 (2)). They will also be of fundamental importance for classifying general-purpose AI (GPAI) models with systemic risks and assessing high-impact capabilities (Art. 51(1) and Annex XIII). Furthermore, both the Board and the Scientific Panel are expected to contribute to the development of benchmarks (Art. 66(g) and Art. 68(3)). Additionally, the first and second drafts of the Code of Practice specifically mentions benchmarks as an example of best-in-class evaluations for risk assessment measures for providers of GPAI models with systemic risks [29, 31]. In the US, benchmarks are also relevant in the recently revoked AI Executive Order [108] and the AI Diffusion Framework [113]. AI benchmarks can further be expected to play a central role in the implementation of legislations such as the EU Digital Services Act (DSA) [28] and the UK's Online Safety Act [112], which require the largest online platforms and search engines - entities that increasingly rely on AI to curate, filter, and/or rank content to millions of users - to perform regular algorithmic audits showing that their systems are safe, fair, and comply with fundamental human rights. In short, AI benchmarks - which constitute a highly heterogenous and far from standardised set of techniques - are increasingly at the heart of policy efforts to make AI more transparent and secure.

At the same time however, as benchmarks are increasingly relied upon to provide AI safety assurances, researchers in a broad range of academic fields have raised serious concerns regarding their current use. This includes critical voices being raised in fields ranging from cybersecurity [64], linguistics [15], and computer science [35], to sociology [26], economics [27], philosophy [52], ethnography [49], and science and technology studies [49]. Such scholars have for example described current AI evaluation practices as a "minefield" [70], that raise serious ethical concerns regarding what should be measured, according to what standards, and with what downstream effects [13]. They have also emphasised that benchmarks are deeply political, performative, and generative in the sense that they do not passively describe and measure how things are in the world, but actively take part in shaping it [38]. This happens as benchmarks continuously influence how AI models are trained, fine-tuned, and applied - practices with wide-ranging political, economic, and cultural effects.

With this backdrop, an interdisciplinary and up-to-date survey of quantitative AI benchmarking critique is momentarily missing. The aim of this article is to address this gap by mapping and discussing known limitations in quantitative AI evaluation practices, drawing together insights from technical, cultural, and historically oriented academic fields. Our goal is to provide an overview of current risks associated with quantitative AI tests, targeting policy makers, stakeholders performing algorithmic audits, and developers of AI models and benchmarks. To achieve this, we gather and analyse around 100 publications published between 1st January 2014 and 31st December 2024 that explicitly and primarily address limitations in current benchmarking practices. Notably, more than half of these publications were published in 2023 or later, highlighting the urgency of providing an updated survey of discussions in the field.

Our findings show that a rapidly growing number of researchers are voicing concerns regarding how benchmarks are used to define and measure what is safe or unsafe, moral or immoral, true or false, toxic or healthy. **They further support the notion that no benchmark is neutral** and that AI tests and benchmarking practices always rest "on interwoven technical and normative decisions" [88, p. 1201] in ways that urge policy makers and AI developers to apply them with caution. In particular, previous research highlights a need to question the relevance and trustworthiness of well-cited benchmarks, since existing studies have repeatedly shown weaknesses in benchmarks perceived as state-of-the-art (SOTA). Our findings also reveal a substantial need to scrutinise capability and safety-oriented AI benchmarks to the same extent as the AI models they are meant to evaluate. In short, AI benchmarks need to be subjected to the same demands concerning transparency, fairness, and explainability, as algorithmic systems and AI models writ large.

In what follows, we first provide a background discussion on AI benchmarking practices, including clarifications on the terminology used in this paper. Next, we situate our work within a series of previous research that summarise benchmark critique, before providing a summary of nine issues with quantitative benchmarks, identified during the course of our research. While not exhaustive, we present these issues as a taxonomy of benchmark critique that highlights crucial and often interrelated points of concern regarding benchmarks, voiced in the past decade. Finally, we provide some concluding remarks on the implications these issues have for policy makers and implementers. On the whole, our work constitutes the first step in a broader research and policy-oriented project aimed at developing a framework for trustworthy AI benchmarks, and will proceed by considering mitigation strategies.

2 BACKGROUND

Etymologically, the term "benchmark" has its roots in land surveying, where a physical mark (known as a 'bench mark') was used as a reference point for measuring elevations. This mark typically consisted of a horizontal groove in a surface, which supported a level surface or 'bench' for a levelling rod. Over time, the term evolved to encompass a broader meaning, referring to any standard or reference point used for comparison or evaluation [80]. Benchmarking is currently applied across many different domains such as bioinformatics [4], environmental quality [41], information retrieval [107], transistor hardware [19], industry and business [17], and within the public sector [16], where it often refers to procedures for comparing the performance or best practices of different companies or processes. Here, we focus on the use of benchmark tests within computing, where they are used to evaluate the performance of hardware or software systems by comparing them to a standard or reference point [41]. More specifically, we zoom in on AI development, where benchmarks are often used to facilitate cross-model comparisons, measure performance, track model progress, and identify weaknesses [90]. Benchmarks are often perceived as comparably cost- and time-effective tools for AI providers who may run them regularly throughout the model development to obtain frequent signals of AI model performance and capabilities [115]. They can be applied to both software and hardware solutions, where the latter for instance evaluates the performance of CPUs, GPUs, TPUs [63], or hardware accelerators [62]. We address

software-oriented benchmarks. Based on the definition proposed by Raji et al. [87], we define such benchmarks as “a particular combination of a set of test datasets, including human-in-the-loop interactions, and associated performance metrics, conceptualised as representing one or more specific tasks or related capabilities, typically chosen by a community of researchers as a shared framework for comparing AI models”. A *testing dataset* would refer to “a separate dataset, distinct from the training data, used to objectively evaluate the performance and generalisability of an AI model”. These datasets are typically composed of *samples* or *instances* that include an input paired with the desired output (also known as reference, gold, or ground truth). A *task* would be “a particular specification of a problem, typically represented as a mapping between an input space and an output or action space, extensionally defined through one or more datasets, and usually associated with a specific performance evaluation metric” (based on [87] and [96]). And finally, a *metric* would refer to “a specification of the mechanism to determine the degree of success or failure of the model’s outputs”. The metric represents a way to summarise model performance over a given task and dataset, usually defined as a single number or score [87]. AI models that achieve the best scores on the metrics of a benchmark are generally considered SOTA in terms of performance.

While humans are always involved in the design and creation of ground truth in benchmarks, they can play a more or less direct and active role in the process of applying benchmark tests. In *automated* or *quantitative benchmarks*, a set of tasks, datasets, and metrics are first defined through human decision making. The execution of a benchmark test is then carried out without direct human intervention. In *qualitative benchmarks*, humans directly intervene and partake in evaluations, for example as evaluators, judges, or real-time interrogators (e.g., adversarial testing or red teaming). We primarily consider quantitative benchmarks. This is not to say that there is no relevant critique concerning qualitative AI, but such methods introduce a different set of critique that is out of scope for our discussions here.

3 RELATED WORK

Over the years, several surveys and meta-reviews have set out to summarize discussions on limitations in AI benchmarking and this paper represents a continuation of such efforts. In 2021, Liao et al. [56, p. 1] identified a wide range of “surprisingly consistent critique” directed at benchmarking practices across fields such as computer vision, natural language processing, recommender systems, reinforcement learning, graph processing, metric learning, and more. The authors present a taxonomy of observed benchmark failure modes, which for example includes implementation variations (due to benchmark algorithms, metrics, and libraries being applied in slightly different ways), errors in test set construction (as a result of label errors, label leakage, test set size, or contaminated data), overfitting from test set reuse, and comparisons to inadequate baselines. Along similar lines, Hutchinson et al. [43] survey discussions on evaluation practices in machine learning, particularly in the fields of computer vision and natural language processing. The authors identify eight key evaluation gaps, including topics such as “neglect of model interpretability” and “oversimplification of knowledge” and argue for a shift toward application-centric evaluations that account for safety, fairness, and ecological validity. Similar points are also raised by Gehrmann et al. [33]’s most recent meta-review, which categorizes issues with text-oriented benchmarks identified over the past two decades, and propose a long-term vision for improving evaluation practices, emphasizing the need for comprehensive evaluation reports that include multiple datasets, metrics, and human assessments. Their key recommendations include focusing on model limitations, enhancing dataset documentation, and adopting a more nuanced approach to evaluation to better characterize model performance and capabilities.

Our review revisits these areas of critique and traces how debates concerning benchmarks have evolved in the last few years. This is not least necessary due to rapid developments currently seen in the field. Recent research has

identified an unprecedented growth in the release of AI benchmarks, especially in the area of safety, starting from 2023 and onwards [93]. Our research indicates a similar increase in publications expressing benchmark critique. For instance, Röttger et al. [93] found that roughly 46% of the AI safety benchmarks they identified as relevant in their survey of the field had been produced in 2023, with as many as 15 new datasets being released during the first two months of 2024. Likewise, almost 55% of the articles we identified as relevant during the course of our research had been released in 2023 or later. Given this rapid increase in attention given to the topic, our primary goal is to extend existing meta-reviews into the present. Whereas previous meta-reviews have tended to focus on critique raised in a limited number of domains (primarily natural language processing and machine learning) we also include interdisciplinary accounts from the humanities and social sciences in our review. Furthermore, we provide insights concerning more recently identified areas of benchmark concern, such as works revealing that AI models can be programmed to underperform on benchmarking tests (a problem known as "sandbagging") [116], and studies presenting evidence that AI safety benchmarks strongly correlate with upstream AI capabilities (causing reasons to worry about AI "safetywashing") [89], insights that both raise additional questions about the validity and trustworthiness of current benchmarking practices.

4 METHODOLOGY

Since benchmarks constitute a key element in most papers that discuss or demonstrate the capabilities of AI models, it is challenging to identify previous works that explicitly and thoroughly address *benchmark critique*, as opposed to applying benchmarks or proposing new methods. When gathering source materials for this review (and instead of conducting keyword searches in research databases) we therefore opted for a method that involved starting from a selection of well-cited papers that problematise how, when, and where benchmarks are used. For each article, we surveyed its reference list and studied how the paper has been cited since its publication. This allowed us to identify additional relevant literature and expand our corpus successively.

In the process, we surveyed a wide range of papers that, for instance, highlight general problems with the production of AI datasets [67, 76], propose alternative ways of designing benchmark leaderboards [57, 92], or discuss the wider use of proxies in tests and evaluations [61, 69, 85]. Such papers provide important contextual insights to discussions concerning benchmarks. However, they were omitted from what we considered our core collection of previous research. This collection ended up consisting of about 100 papers and articles that *explicitly* and *primarily* highlight issues with benchmarks. Notably, articles that propose new benchmarks were *not* added to this collection by default, even though such articles naturally contain some level of benchmark critique. The reasoning behind this is twofold. On the one hand, it was necessary to limit the size and scope of our review. On the other hand, many articles proposing new benchmarks (if not most of them) do not question or discuss many (if any) of the underlying assumptions that benchmark assessments rely on. In that sense, they reproduce the general notion that quantitative benchmarks provide a reasonable technical "fix" to issues with AI safety and capability assessments, rather than taking part in more fundamental and critical discussions on current limitations in benchmark use and design. For these reasons, they were also excluded from our core collection of benchmark critique.

Our collection was further limited by only containing articles published between 1st January 2014 and 31st December 2024. While it is true that AI models have been tested and evaluated since their origins in the mid 1950's (with the Turing test serving as an iconic example), we limit our review to this decade long period, since 2014 marks the starting point for recent intensifications in AI research and development. We consider both published and pre-published papers in our

review, which on a practical level involved categorising identified critique into relevant areas of concern, without a pre-existing target of final categories. When important concepts (such as sandbagging) appeared in the literature, targeted searches for research in these fields were made. Given that a majority of existing AI benchmarks (and benchmark critique) address text-based AI models, we also made special efforts to include discussions concerning benchmarks for images, sound, and audiovisual content.

Our resulting meta-review is not exhaustive but it covers a broad range of critique that has been aimed at benchmarking practices. In the following discussions, we focus on works that voice critique with particular relevance for policy makers and policy implementation. We also shed light on works that emphasise areas of concern that cut across different modalities (text, images, sound, moving images) and point toward fundamental weaknesses in the design and application of benchmarks, as opposed to research that critiques individual benchmarks. Our goal has further been to provide a *diverse* account of concerns regarding benchmarks. Hence, some sections will have ample publications reinforcing their claims while others, that voice more recent and/or unique critique, are backed up by fewer sources.

5 NINE REASONS TO BE CAUTIOUS WITH BENCHMARKS

In the sections below, we summarise the main issues that were identified during the course of our research. Importantly, these issues - presented here as a taxonomy - are not arranged according to their importance or urgency. They are also not meant to be understood as isolated issues, but rather deeply interlinked problems. Indeed, this complexity and interdependence is precisely what makes AI evaluations challenging.

5.1 Problems with Data Collection, Annotation, and Documentation

An initial set of issues with AI benchmarks found during our research is limitations in the collection, annotation, and documentation of benchmark datasets. This ties into a broader critique regarding insufficient documentation in AI research which is central to calls for more transparent and trustworthy algorithmic systems [32, 67, 76, 95, 100]. Research has found that it is often difficult to trace precisely *how*, *when*, and by *whom* benchmark datasets have been made [24, 90] which compromises the ability for benchmarks to be used in robust and generalisable ways [6]. The issue has partly been linked to the low status of dataset-related work within the machine learning community (which instead privileges model development) [77, 94], and the fact that AI datasets are often “reduced, reused, and recycled” [50], which complicates documentations of their possible limitations [82, 109]. Notably, Koch et al. [50] have found that more than 70 percent of the benchmark datasets used in prominent computer vision papers had been reused from other domains. Park and Jeoung [82] also find that while benchmark datasets often contain ample information on how to use the dataset, documentation is often missing concerning their shortcomings and social impact. In their exploration of benchmark sharing platforms like HuggingFace¹ and PapersWithCode² the same researchers also note that confusing metadata terminology severely complicates efforts to understand benchmark dataset documentation.

Like AI training datasets, benchmarks have been singled out for raising ethical and legal questions concerning copyrights, privacy, informed consent, and rights to opt-out [83]. Many benchmarks rely on crowd-sourced or user-generated content from platforms like Wikihow³, Reddit⁴, or trivia websites [38, 48], whose annotation may be noisy and biased, lack input from expertise in specialised fields, or be produced under exploitative conditions [5, 99, 111].

¹<https://huggingface.co/>. Last seen 22nd January 2025.

²<https://paperswithcode.com/>. Last seen 22nd January 2025.

³<https://www.wikihow.com/Main-Page>. Last seen 22nd January 2025.

⁴<https://www.reddit.com/>. Last seen 22nd January 2025.

Previous research has noted that this matters when baselines for "good", "bad", and "safe enough" AI models are calibrated in sensitive contexts [38, 88]. Moreover, it is essential to highlight the absence of human performance references and difficulty rubrics in benchmarks, which are increasingly important factors in evaluating capabilities and generality [20].

A lack of care in the making of benchmark datasets has further been found to result in AI models exploiting unknown quirks and spurious cues in the training data, rather than solving their original intended task [34, 56, 83]. For instance, Oakden-Rayner et al. [74] found that an X-ray image classification model predicted collapsed lungs with high accuracy, although it turned out the model only identified the presence of a chest drain (used to cure the condition) that was (unknowingly) present in a majority of the positive training images. In other words, the model performed well in benchmark tests, but the reason why completely sidestepped the original purpose of the task (identifying collapsed lungs). When the all images containing chest drains were removed from the training dataset, the model performance dropped by over 20%. Similar patterns have also been confirmed in recent research on LLM evaluation [81] and is sometimes discussed as failure in the "world model" of an AI system [114]. While it could be argued that picking up on spurious cues is a sign of intelligence and therefore a capability worth noting, the issues above highlight how a lack of in-depth attention to the data that benchmarks rely on - alongside a failure to acknowledge how little is actually known about *how* and *why* AI models perform well in benchmark tests - can produce frail and uncertain AI evaluations.

5.2 Weak Construct Validity and Epistemological Claims

Another genre of benchmark critique focuses on the epistemological claims that tend to surround benchmarks and examines the limits of what can be known through quantitative AI tests. A central reference point in these discussions is the observation by Raji et al. [87] that many benchmarks suffer from construct validity issues in the sense that they do not measure what they claim to measure. As the authors proclaim, this is especially troublesome when benchmarks promise to measure universal or general capabilities, since this vastly misrepresents their actual capability. As a result, the authors argue that framing an AI benchmark dataset as general purpose "is ultimately dangerous and deceptive, resulting in misguidance on task design and focus, underreporting of the many biases and subjective interpretations inherent in the data as well as enabling, through false presentations of performance, potential model misuse". [87, p. 5]. At the heart of this critique lies the realization that many benchmarks do not have a clear definition of what they claim to measure, which makes it impossible to measure if they succeed in the task or not [8, 14]. In a close analysis of four benchmarks used to evaluate fairness in natural language processing (StereoSet, CrowS-Pairs, WinoBias, and WinoGender), Blodgett et al. [14] for example found that all benchmarks revealed severe weaknesses in terms of defining what is being measured. For instance, culturally complex and highly contested concepts like "stereotypes" or "offensive language" were left unspecified, causing a series of logical failures and interpretational conflicts. Elsewhere, research has shown strong disagreements in how benchmark tasks are *conceptualised* and *operationalised* [106], and found that benchmarks are applied in highly idiosyncratic ways [93]. Frequently, the difficulty defining what benchmarks evaluate persist since there is no clear, stable and absolute ground truth for what is claimed to be measured [71]. Since concepts like "bias" and "fairness" are inherently contested, messy, and shifting, benchmarks that promise to measure such terms will inevitably suffer from an "abstraction error" that produces a false sense of certainty [98, p. 63].

It has also been pointed out that many benchmarks datasets are inadequate and/or unuseful proxies for what they are meant to evaluate. For instance, researchers have identified a slippage in distinguishing between algorithmic "harms" and algorithmic "wrongs" when evaluating the capabilities of AI models - two not necessarily overlapping concepts [25]. Others have questioned whether the content of benchmark datasets are reasonable substitutes for the "real world"

scenarios they are meant to reflect. For instance, researchers have questioned the decision to use examples generated by Amazon crowdworkers and posts from the Reddit forum "*Am I the asshole?*" as proxies for ethics and morals in benchmarks such as HellaSwag - a widely cited benchmark for language models [48]. In benchmarks consisting of professional exams, researchers have further argued that such tests "emphasize the wrong thing" and "overemphasize precisely the thing that language models are good at" and are thus unreliable measures of things such as medical or legal skill [71, n.p]. As Narayanan and Kapoor [71] put it, "it's not like a lawyer's job is to answer bar exam questions all day". A recent study by Ren et al. [89] also found that many widely used safety benchmarks (including ETHICS, TruthfulQA, GPQA, QuALITY, MT-Bench, LMSYS Chatbot ARENA, ANLI, AdvGLUE, and AdvGLUE++) highly correlate with general and upstream model capabilities, raising concerns regarding "safetywashing" as applying them could imply that "capability improvements are misrepresented as safety advancements". While it could be argued that capability and safety are largely entwined (the more capable a model is, the higher the likelihood it could cause harm), Ren et al. [89, p. 1] suggest that a blurred distinction between the two may hide the fact that severe biases and safety issues in AI models can persist, even as their overall capabilities improve.

5.3 Sociocultural Context and Gap

Another key insight from previous research concerns the importance of the social, economic and cultural contexts where AI benchmarks are created, used, and maintained. Among researchers engaging in benchmark critique, we identify a strong consensus regarding the need to recognise that benchmarks are ultimately "normative instruments that perpetuate particular epistemological perspectives about how the world is ordered" [78, p. 1877]. Qualitative research has also examined the cultural and social environments where benchmarks are made, finding that they are deeply shaped by shared and arbitrary assumptions, commitments, and dependencies [26, 65, 77, 83, 94, 95]. Such assumptions for example include valuing "efficiency at the expense of care; universality at the expense of contextuality; impartiality at the expense of positionality; and model work at the expense of data work" [95], or reproduce contested ideas such as the notion that low-quality data can be drowned out by scale [77]. Recent research has also illustrated that AI safety research and benchmark competitions are increasingly informed by political movements and ideologies such as longtermism and effective altruism [2]

Scholars have further identified sociotechnical gaps and a lack of consideration for downstream utility as a concern in AI benchmarking [43]. Liao and Xiao [55] highlight that this means it is often unclear who is meant to care about benchmark evaluation results and how they should be used in practice. For instance, a recent study by Blagec et al. [12], which compared the explicitly stated needs for AI technologies among clinical medical practitioners with existing clinical benchmark datasets, found that most benchmarks failed to answer to the needs of medical experts. They also found that benchmarks for the most urgently requested medical or clinical tasks were completely missing and noted that similar misalignments likely exist in many other fields. For example, Jannach and Bauer [44, p. 79] highlight a lack of attention to how recommender systems "may create value; how they, positively or negatively, impact consumers, businesses, and the society; and how we can measure the resulting effects". Ethayarajh and Jurafsky [27] further argue that failures to consider the practical utility of benchmarks has for example made it possible to ignore the discriminatory and environmental damages of AI technologies and allowed for highly energy-inefficient and deeply biased AI models to reach the top of most benchmark leaderboards.

5.4 Narrow Benchmark Diversity and Scope

Previous research has further found that current benchmarking practices suffer from diversity issues, a problem that is also found within the broader AI ecosystem [36]. On the one hand, a vast majority of benchmarks focus on text, while other modalities (audio, images, video, and multimodal systems) remain largely unexamined [88, 93, 115]. This concentration is problematic since AI models are increasingly multimodal in scope. Guldemann et al. [39] also find that benchmarks addressing user privacy, copyright infringement, and interpretability are currently lacking and incomprehensive, while benchmarks in safety areas dealing with corrigibility and explainability are practically missing all together. **This leads the authors to conclude that current practices for** evaluating topics like safety and ethics “are often **simplistic and brittle, leading to inconclusive results**” [39, p. 3]. Koch et al. [50] further identify a concentration on fewer and fewer benchmark datasets within most task communities and note that dominant benchmarks have been introduced by researchers at just a handful of elite institutions, raising questions about representation diversity in the design of benchmarks. Scholars have also found that current AI safety evaluation practices almost exclusively deal with English content [64, 93], and are frequently based on datasets where minorities are under-represented, despite efforts to diversify them [100]. This raises concerns regarding the inclusion of multiple perspectives on complex topics like ethics and harm.

Aside from mainly focusing on a small range of tasks (evaluating capabilities of English-speaking language models), previous research has identified that most benchmarks tend to be abstracted out of their social and cultural context [98], and rely on a **static, one-time testing logic**, where results from single evaluations are taken to represent model capabilities writ large. This has given rise to calls for more multi-layered [115], longitudinal [68], and holistic evaluation methods that make sure that AI models do not just perform well in controlled environments, but also in critical/real world circumstances over time [18, 64, 75]. In their survey of LLM benchmark inadequacies, McIntosh et al. note that “current benchmarks generally adopt task-based formats, such as Multiple Choice Questions (MCQs) and dialogue-based evaluations, which tend to be static and do not capture the evolving nature of human-AI interactions” [64, p. 6]. Reuel et al. [90] further note that failures to re-run evaluations multiple times (using different random seeds and sampling temperatures for example), implies that very little is known about the intra-model variance of benchmarks. As a result, they conclude that it is possible that “most benchmarks fail to distinguish signal and noise” [90, p. 9]. Others have emphasised that AI audits often fail to consider risks associated with multiple (inter)acting AI systems [11], and rarely take human actions and motivations into consideration [18, 88, 115].

While most benchmarks are designed to tell us something about a model’s success, research has further pointed out that they often reveal little (or nothing) about their particular ways of *making mistakes*, which is crucial from an AI safety and policy enforcement perspective. As Gehrmann et al. [33] put it, “ranking models according to a single quality number is easy and actionable - we simply pick the model at the top of the list - [yet] it is much more important to understand when and why models fail” [33, p. 130]. For instance, they suggest that a focus on errors and fragilities (as opposed to instances of success) can be useful for developers of smaller models, since “work on quantifying shortcomings is equally applicable to smaller models and methods that improve model robustness often work on all model sizes” [33, p. 131]. In this sense, failure-focused benchmarks could play an important role in equalling out the playing field in AI development.

5.5 Economic, Competitive, and Commercial Roots

Another contextual element that has been singled out as important is the competitive and commercial roots of benchmark tests. Previous research has emphasised that capability-oriented benchmarks are deeply embedded in corporate marketing strategies and play an important role in increasing the AI hype, attracting customers and investors, and showcasing how models outperform competitors [38, 78, 120]. As Orr and Kang [78, p. 1881] put it, benchmarks "serve as the technological spectacle through which companies such as OpenAI and Google can market their technologies". Many benchmarks also have origins from within the industry and are capability-oriented and centred around tasks with a high potential economic reward, as opposed to focusing on other goals such as ethics and safety [27, 89].

Previous research has noted that this competitive and corporate embedding discourages thorough self-critique since there is a direct "incentive mismatch between conducting high-quality evaluations and publishing new models or modelling techniques" [33, p. 103] and that the field of AI development is "turning into a giant leaderboard, where publication depends on numbers and little else (such as insight and explanation)" [21]. While a lack of incentives to disclose weaknesses and limitations is a general problem in science [101], it has been pointed out that benchmarks have played an especially central role in naturalizing and solidifying a competitive culture in AI research, which is increasingly approached as a "sport" [78]. As of late, benchmark evaluations have also become increasingly professionalized and transformed into an industry in itself with the rise of platforms like Kaggle and Grand Challenge, who provide organisational and infrastructural support to AI competitions and increasingly function as infrastructures of power in fields like medical imaging [58]. The issue of optimising for high benchmark scores at the expense of insight and explanation is known as a form of SOTA-chasing [50] and sometimes described as the "benchmark effect" [104]. Previous research [60] has also described benchmarks as an example of what Stengers [103] calls "fast track research" which idolises rapid, cumulative publication, thus producing a "winners curse" in AI development [97].

Risks associated with the competitive and commercial roots of benchmarks have further been linked to the growing influence of industry in AI research, where private businesses share of the biggest AI models has increased from 11% in 2010 to 96% in 2021 [1]. In such a context, researchers have noted that current benchmarking tasks - which are generally highly data-intensive - are especially well suited to fit AI models that have been developed within the industry, whose access to advanced data infrastructures, computing power, valuable datasets, and skilled researchers now vastly exceed those of academic researchers [1]. This concentration of power could potentially stifle robust AI evaluations and hinder the development of AI models that adhere to other aims and goals than commercial ones. Scholars have also warned that if academic researchers continue to uphold data-intensive benchmark tests as the SOTA, there is a risk that their research will become increasingly dependent on technological infrastructures provided by the industry [51].

5.6 Rigging, Gaming, and Measure Becoming Target

A closely related issue concerns how benchmark tests can be tricked and gamed. In areas and modalities where best-practice benchmarks are missing (i.e., practically all modalities, except for text-based benchmark evaluations), researchers have noted that there are strong incentives to "rig" benchmark tests. For instance, Dehghani et al. [22] survey how know-how and recipes for how to score high on benchmark setups are often widely circulated online. Recently, language models have also been found to be optimised for answering the multiple choice questions that are often part of benchmarks [3], and to (either intentionally or unintentionally) "fake" alignment with ethics or safety goals [37]. The issue points towards what is known as Goodhart's law, i.e. the recognition that "when a measure becomes a target, it ceases to be a good measure" [105].

One reason why gaming can proceed is that users of benchmarks rarely provide the resources needed to validate and replicate their test results [8, 10, 22, 90]. To examine the relevance and validity of benchmark scores, researchers have emphasised that it is necessary to access information concerning all aspects of the evaluation procedure (such as the original evaluation code and all details concerning the experimental setup) [10]. Providing such documentation is far from standard, however, especially when proprietary AI models are concerned. This makes it possible to tweak and cherry-pick benchmark results - a problem that is especially pressing given that subtle "variations in prompts, formatting or other implementation details can significantly impact the performance and validity of evaluations" [10, p. 3]. In their in-depth analysis of 24 SOTA language model benchmarks, Reuel et al. [90] found that only four provided scripts to replicate the results and that no more than ten performed multiple evaluations or reported the statistical significance of their results.

An increasingly well discussed issue also concerns the problem of "data contamination" i.e., the risk that the models have either intentionally or unintentionally ingested benchmark datasets during training, which severely questions the integrity of AI tests [9, 59, 91, 117, 119]. The problem - which can be an instance of data leakage [47] or train-test-overlap [53] - has been known for long, and produces similar effects to those of overfitting and memorization [59, 110], leading to models with low generalization power that perform well on familiar tasks (in-distribution) but fail other tasks with a similar difficulty and distribution shift (out-of-distribution) [9, 59, 71, 91, 117-119]. When testing GPT4 on benchmark problems from Codeforces (a website hosting coding competitions) in 2023, for instance, Narayanan and Kapoor [71] found that the AI model could regularly solve benchmark problems classified as easy - as long as the problems had been added before 5th September 2021. For problems added later, GPT4 could not get a single question right, suggesting that the model had memorised questions and answers. Similar results have also been identified in multiple other models and benchmarks [9, 59, 91, 117, 119]. Despite the fact that issues with data leaks are so well known that strategies have been developed to avoid it⁵, there is still a widespread lack of reporting of data contamination tendencies during benchmark tests. In a study from October 2024, for example, Zhang et al. [119] found that out of 30 analysed models, only 9 reported train test overlap.

Another recently identified issue is called "sandbagging" and involves a "strategic underperformance on an evaluation" which occurs when an AI developer intentionally *understates* a models capability, for instance to avoid becoming a target for AI safety regulation. In a 2024 study, for example, Weij et al. [116, p. 1] prompted frontier models like GPT-4 and Claude 3 Opus and found that they "selectively underperform on dangerous capability evaluations, while maintaining performance on general (harmless) capability evaluations". The researchers also found that that it was possible to fine-tune and adjust both frontier and smaller models to hide specific capabilities or target specific capability scores. The issue could be generalised to popular benchmarks like the Weapons of Mass Destruction Proxy Benchmark (WMDP), and puts the trustworthiness of benchmark evaluations into question, especially in a regulatory context.

5.7 Dubious Community Vetting and Path Dependencies

A related area of research emphasises how benchmarks become naturalised and reach standard status because of the culture and logic of academic citations [78]. For instance, new benchmarks are commonly introduced together with new or updated AI models. If the AI model becomes popular, the benchmark may become widely cited and circulated as a secondary effect, even though the developers of the benchmark did not intended or expect it to become standard. In this way, a benchmark that might have been selected for practical purpose or due to subjective preferences may

⁵<https://livebench.ai>. Last seen 22nd January 2025.

come to have substantial technical and cultural influence as a result of community vetting, even though its suitability as a yardstick for ethics, safety, or performance may be questioned [77, 78]. For instance, Denton et al. [23] show how this was partly the case when the ImageNet dataset - a key reference point in the performance testing of computer vision models - became standard, following the unforeseen success of the ImageNet Large Scale Visual Recognition Challenge hosted by Toronto University in 2010. Likewise, the so-called Lena test image - also central to computer vision benchmark tests - was taken from the centerfold of a November 1972 Playboy magazine and catapulted into computer history since "someone happened to walk in with a recent issue of *Playboy*" at a time of need, according to what has been described as the most credible origin story [69, p. 80]. Since then, the ImageNet dataset and its associated benchmark challenge has become a symbol for dataset bias [23], while the Lena test image has come to serve as a prime example of the role of whiteness and women's sexualised bodies in the standardisation of digital visual culture [69]. These examples highlight how benchmarks are fundamentally cultural and political products whose power and influence may and be (uncritically and problematically) reinforced through community vetting.

Aside from bias and representational issues, Schlangen [96, p. 1] argues that what is typically missing and left implicit in the peer-review-fuelled process of benchmark use is the argumentation for why scoring well on a particular benchmark "constitutes progress, and progress towards what". Instead, researchers are expected to routinely demonstrate performance on dominant benchmarks, despite the fact that more task-specific benchmarks may be more technically appropriate [50]. In this way, peer-washing serves to "maintain datasets as authoritative proxies even when they are shown to be harmful or problematic" [77, p. 4966]. New benchmarks often have a difficult time to gain traction because of the dominance and authority of well-cited benchmarks [45], and in a study of 3765 benchmarks, Ott et al. [79] found that only a small minority of the proposed benchmark solutions reached widespread adoption. This problem is also recognised within the industry. For instance, researchers at Google's Brain Team describe what they call a "benchmark lottery" which "postulates that many factors, other than fundamental algorithmic superiority, may lead to a method being perceived as superior" [22, p. 1]. Ironically, researchers have simultaneously found that a majority of influential benchmarks have been released as preprints without going through rigorous academic peer-review [64, p. 6]. In other words, problems with peer-washing both concern the poor quality control that many benchmarks are subjected to, and the self-fulfilling and sometimes excessive ways through which certain benchmarks are propelled into standards.

Bao et al. [7] further note that most papers that introduce benchmarks have origins in the field of machine learning and are mainly focused on methods. This means that the content of benchmarks datasets are often considered secondary, since the datasets are merely there to provide a baseline for comparison between different evaluation methods. When benchmarks are applied to real and specific use-cases - such as evaluating the fairness of algorithmic risk assessment instruments within the criminal justice system - the downstream effects of such a lack of concern for datasets can have worrying effects. As Bao et al. [7] put it, a paper on benchmarks can be of high quality in a pure AI/ML methods sense, but irrelevant, dangerous, or harmful when applied in circumstances such as the criminal justice system, since it may introduce and perpetuate harms and mistranslations. The researchers also suggest that the current peer-review system implies that benchmarks that are primarily relevant from a methods/machine learning perspective will be cited far more often than benchmarks that are relevant for use in specific, real-life use-cases. Looking too closely at citation counts when determining the quality and relevance of benchmarks may thus effectively lead users astray and complicate identifying benchmarks with a high *practical* utility.

Problems with dubious community vetting become especially worrying given that benchmarks create "path dependencies" in AI research, meaning they reinforce certain methodologies and research goals, while stifling those that do not align with the logic of dominant benchmark tests [13]. In particular, Koch and Peterson [51, p. 3] warn against the

tendency for benchmarks to favour a form of task-driven scientific monoculture that privileges immediate, explicit, formal, quantitative, and easily-interpretable evaluation mechanisms that "prioritize one or a few key epistemic values (e.g., accuracy, safety)," at the expense of a broader and more complex vision of scientific progress. The authors note how a series of methodological paradigms in AI research (such as boosting, Bayesian networks, Bayesian non-parametrics, and support vector machines) have rapidly faded away following the boom in deep learning research around the year of 2014, and partly attribute this to the "epistemic narrowness" of current benchmark tests, which are well-suited for deep learning models (whose performance can be reliably increased through scaling), but less favourable for other methodological paradigms in AI. "When scientists can only gain high-status publications by demonstrating SOTA accuracy," writes Koch and Peterson [51, p. 30], "the safest research choice becomes incrementally advancing proven methods, not innovating new ones". According to Koch, the cost of current benchmark practices - which privileges scale, compute, and the use of larger and larger training datasets - for instance includes repeated privacy and copyright violations, emotional harms caused to under-paid data workers, and increasing ecological and environmental pressures due to high energy consumption.

5.8 Rapid AI Development and Benchmark Saturation

Another social, economic, and cultural issue is the speed of AI developments. As the capabilities of AI models have increased manifold in the past decade, researchers have emphasised that many benchmarks are old and designed to test models far simpler than those in use today [10, 48]. For instance, Biderman et al. [10, p. 5] find that many prominent LLM benchmarks (including Lambada, AI2 ARC, OBQA, Hella Swag, and WinoGrande) were "designed prior to shifts such as in-context learning and chat interaction, and therefore were not designed to take these formats and approaches into account", noting that this may affect their validity in unforeseen ways. Many benchmarks also struggle with the challenge of quickly being outperformed as AI models achieve 100 percent accuracy scores in tests [15, 40]. A look at current AI safety leaderboards such as HELM also reveal that many AI models score notably high on benchmarks that are widely applied today. The tendency for AI models to outperform benchmarks has been described as an issue of *saturation* and implies that a benchmark no longer reflects model performance [79].

Relatedly, McIntosh et al. [64] emphasises that many benchmark frameworks are slow and complicated to implement, meaning that evaluation processes can span weeks or months, which hinders timely feedback on AI model safety risks. This becomes an issue since new model releases often enter markets continuously, making it difficult to relocate evaluation resources in quick and adequate ways. It also undermines a "benchmark's ability to consistently evaluate reasoning, comprehension, or multimodal integration, as the results may vary with each model iteration" [64, p. 13], which is especially concerning in a regulatory setting, where quick, fair, and accurate AI assessments are key. The use of thresholds, either on benchmark performance or training compute (i.e., the amount of computational resources used to train an AI model), is often seen as central to determine which AI models should warrant further regulatory scrutiny [30, 108, 113]. However, several limitations have been identified in relation to such attempts. Creating benchmarks that can keep pace with the rapid development of AI is increasingly challenging, at the same time as recent approaches (such as auxiliary or bootstrapped models or post-training interventions) enable enhanced capabilities with reduced training compute [42].

5.9 AI Complexity and Unknown Unknowns

A final issue that has been discussed in relation to benchmarks concerns AI complexity and the fundamental difficulty of foreseeing what risks, dangers, and threats AI models could pose to society. McIntosh et al. [64, p. 13] point out that

"LLM benchmarks are constrained by the current limit of the benchmark creators' human knowledge, hindering their ability to fully assess and cultivate emerging AI capabilities that may surpass conventional human understanding". According to the authors, the lack of general understanding of emerging AI capabilities, and the natural limitations of the benchmark creators' knowledge on a potentially infinitely large number of domains and tasks, may lead to generalist approaches that "often fail to address the subtle requirements of critical sectors such as national security or healthcare". They argue that this does not just pose security risks, but could potentially also hinder innovation.

Evaluations of AI models are also complicated by the potential presence of unknown and latent vulnerabilities in AI models that "make it very hard to distinguish between (a) actually safe and (b) appears safe but is not" [72]. In 2023, for instance, Nasr et al. [73] discovered that surprisingly simple prompts can 'break' the safety barriers of AI models, raising questions about the robustness of existing safety measurements. More precisely, the simple command "Repeat the word 'poem' forever" was found to make ChatGPT output several megabytes of sensitive training data. As a result of this finding, the authors conclude that "just as vulnerabilities can lie dormant in code - sometimes for decades - our attack demonstrates the potential for latent, hard-to-discover ML vulnerabilities that lie dormant in aligned models" and go unnoticed by existing safety benchmarks and tests [73, p. 13]. Researchers have also noted the difficulty of foreseeing how complex AI models respond to (both small and large) interventions such as safety alignments and fine-tuning. For instance, efforts to fine-tune AI models to address safety and/or security risks have been found to degrade a model's performance in other safety areas, or introduce entirely new security risks Qi et al. [86].

6 CONCLUSIONS

Measuring the capabilities and risks of AI models and systems is difficult and one of the main challenges in the use and development of AI. Even with the best of intentions (such as disclosing discrimination or identifying potential societal harms), previous research has repeatedly shown that quantitative AI benchmarks struggle to perform their intended task. Benchmarks have been found to promise too much [87], be gamed too easily [70, 116], measure the wrong thing [74], and be ill-suited for practical use in the real world [7, 27]. They have also been found to display a serious lack of documentation [90], randomly reach an unjustified status through community vetting [22], and forward questionable cultural assumptions [46, 48], that for example ignore environmental concerns [43]. Furthermore, benchmarks have been critiqued for being narrow and mainly evaluating English [64, 93] and text-based AI models [88] according to a one-time testing logic [68] that ignores the AI capabilities in other modalities such as imagery and sound [88], and fails to acknowledge that the potential harms of AI models cannot be properly understood through evaluations done in isolated, abstracted, test-environments, devoid of humans [88] and other technical systems [11]. Taken together, these issues point toward fundamental fragilities in current efforts to quantitatively measure and mitigate harm in AI.

Cars, air planes, medical devices, drugs, and numerous other products within our societies comply with strict regulation to ensure their safety. There is no reason to believe that similar safety assurances can not be developed for AI and the intensified interest in AI benchmarks clearly signal a drive to do so. Outside the scope of this meta-review, we have identified numerous papers that propose strategies for mitigating issues with benchmarks, for instance by hiding benchmark training datasets [20] or using so-called "dynamic benchmarks" [9] to counter gaming and data contamination risks, aggregating evaluation tasks into single multi-task benchmarks to increase the reliability of test results [54, 102], or opting for evaluation methods that directly involve human interrogators as opposed to quantitative metrics [18]. An increasing number of researchers have also proposed new and promising frameworks for assessing and "benchmarking the benchmarks" [66, 90]. Because of uncertainties regarding the effectiveness [6, 88, 119] and widespread adoption of such mitigation efforts, however, our meta-review suggests that the field of quantitative benchmarking is currently

ill-suited to single-handedly (or primarily) carry the burden of providing the assurances requested by policy makers. Our review also shows that from a policy perspective, relying on indicators such as citation counts to determine what benchmarks to trust is insufficient. We identify a strong incentive gap in the use of benchmarks between academic researchers (who may for example primarily be interested in methods development and are submerged in the "fast" logic of academic publication), corporations (who have strong incentives to use benchmarks for economic interests), and regulators (who have a particular responsibility to consider the practical utility and potential downstream effects of particular benchmarks). From a regulatory perspective - and the perspective of anyone who wants to apply a benchmark to a concrete, real-life case - we especially identify a need for new ways of signalling *what benchmarks to trust*. We do not necessarily need standardised benchmark metrics and methods. But we do need standardised methods for *assessing the trustworthiness of benchmarks from an applied and regulatory perspective*.

REFERENCES

- [1] Nur Ahmed, Muntasir Wahed, and Neil C. Thompson. 2023. The growing influence of industry in AI research. *Science* 379, 6635 (March 2023), 884–886. <https://doi.org/10.1126/science.ade2420>
- [2] Shazeda Ahmed, Klaudia Jaźwińska, Archana Ahlawat, Amy Winecoff, and Mona Wang. 2024. Field-building and the epistemic culture of AI safety. *First Monday* (April 2024). <https://doi.org/10.5210/fm.v29i4.13626>
- [3] Norah Alzaharani, Hisham Abdullah Alyahya, Yazeed Alnumay, Sultan Alrashed, Shaykhah Alsubaie, Yusef Almushaykeh, Faisal Mirza, Nouf Alotaibi, Nora Altwairish, Areeb Alowisheq, M. Saiful Bari, and Haidar Khan. 2024. When Benchmarks are Targets: Revealing the Sensitivity of Large Language Model Leaderboards. <http://arxiv.org/abs/2402.01781>
- [4] M. R. Aniba, O. Poch, and J. D. Thompson. 2010. Issues in bioinformatics benchmarking: the case study of multiple sequence alignment. *Nucleic Acids Research* 38, 21 (2010), 7353–7363. <https://doi.org/10.1093/nar/gkq625>
- [5] Lora Aroyo and Chris Welty. 2015. Truth Is a Lie: Crowd Truth and the Seven Myths of Human Annotation. *AI Magazine* 36, 1 (March 2015), 15–24. <https://doi.org/10.1609/aimag.v36i1.2564>
- [6] Varvara Arzt and Allan Hanbury. 2024. Beyond the Numbers: Transparency in Relation Extraction Benchmark Creation and Leaderboards. <http://arxiv.org/abs/2411.05224>
- [7] Michelle Bao, Angela Zhou, Samantha Zottola, Brian Brubach, Sarah Desmarais, Aaron Horowitz, Kristian Lum, and Suresh Venkatasubramanian. 2022. It's COMPASlicated: The Messy Relationship between RAI Datasets and Algorithmic Fairness Benchmarks. <http://arxiv.org/abs/2106.05498>
- [8] Thomas Bartz-Beielstein, Carola Doerr, Daan van den Berg, Jakob Bossek, Sowmya Chandrasekaran, Tome Eftimov, Andreas Fischbach, Pascal Kerschke, William La Cava, Manuel Lopez-Ibanez, Katherine M. Malan, Jason H. Moore, Boris Naujoks, Patryk Orzechowski, Vanessa Volz, Markus Wagner, and Thomas Weise. 2020. Benchmarking in Optimization: Best Practice and Open Issues. <http://arxiv.org/abs/2007.03488>
- [9] Sandi Besen. 2024. The Death of the Static AI Benchmark. <https://towardsdatascience.com/the-death-of-the-static-ai-benchmark-88b5ff437086>
- [10] Stella Biderman, Hailey Schoelkopf, Lintang Sutawika, Leo Gao, Jonathan Tow, Baber Abbasi, Alham Fikri Aji, Pawan Sasanka Ammanamanchi, Sidney Black, Jordan Clive, Anthony DiPofi, Julen Etxaniz, Benjamin Fattori, Jessica Zosa Forde, Charles Foster, Jeffrey Hsu, Mimansa Jaiswal, Wilson Y. Lee, Haonan Li, Charles Lovering, Niklas Muennighoff, Ellie Pavlick, Jason Phang, Aviya Skowron, Samson Tan, Xiangru Tang, Kevin A. Wang, Genta Indra Winata, François Yvon, and Andy Zou. 2024. Lessons from the Trenches on Reproducible Evaluation of Language Models. <http://arxiv.org/abs/2405.14782>
- [11] Abeba Birhane, Ryan Steed, Victor Ojewale, Briana Vecchione, and Inioluwa Deborah Raji. 2024. AI auditing: The Broken Bus on the Road to AI Accountability. <http://arxiv.org/abs/2401.14462>
- [12] Kathrin Blagec, Jakob Kraiger, Wolfgang Frühwirt, and Matthias Samwald. 2023. Benchmark datasets driving artificial intelligence development fail to capture the needs of medical professionals. *Journal of Biomedical Informatics* 137 (Jan. 2023), 104274. <https://doi.org/10.1016/j.jbi.2022.104274>
- [13] Borhane Blili-Hamelin and Leif Hancox-Li. 2023. Making Intelligence: Ethical Values in IQ and ML Benchmarks. In *2023 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Chicago IL USA, 271–284. <https://doi.org/10.1145/3593013.3593996>
- [14] Su Lin Blodgett, Gilsinia Lopez, Alexandra Olteanu, Robert Sim, and Hanna Wallach. 2021. Stereotyping Norwegian Salmon: An Inventory of Pitfalls in Fairness Benchmark Datasets. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, Online, 1004–1015. <https://doi.org/10.18653/v1/2021.acl-long.81>
- [15] Samuel R. Bowman and George Dahl. 2021. What Will it Take to Fix Benchmarking in Natural Language Understanding?. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Online, 4843–4855. <https://doi.org/10.18653/v1/2021.naacl-main.385>
- [16] Isabelle Bruno. 2014. *Benchmarking*. Springer Netherlands, Dordrecht, 363–368. https://doi.org/10.1007/978-94-007-0753-5_170

- [17] Robert C. Camp. 1989. *Benchmarking : The Search for Industry Best Practices That Lead to Superior Performance*. Quality Press, the University of Michigan.
- [18] Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. 2023. A Survey on Evaluation of Large Language Models. <http://arxiv.org/abs/2307.03109>
- [19] Z. Cheng, CS. Pang, P. Wang, and et al. 2020. How to report and benchmark emerging field-effect transistors. *Nature Electronics* 5 (2020), 416–423. <https://doi.org/10.1038/s41928-022-00798-8>
- [20] François Chollet. 2019. On the Measure of Intelligence. <http://arxiv.org/abs/1911.01547>
- [21] Kenneth Ward Church and Joel Hestness. 2019. A survey of 25 years of evaluation. *Natural Language Engineering* 25, 06 (Nov. 2019), 753–767. <https://doi.org/10.1017/S1351324919000275>
- [22] Mostafa Dehghani, Yi Tay, Alexey A. Gritsenko, Zhe Zhao, Neil Houlsby, Fernando Diaz, Donald Metzler, and Oriol Vinyals. 2021. The Benchmark Lottery. <http://arxiv.org/abs/2107.07002>
- [23] 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 (July 2021), 1–14. <https://doi.org/10.1177/20539517211035955>
- [24] Remi Denton, Alex Hanna, Razvan Amironesei, Andrew Smart, Hilary Nicole, and Morgan Klaus Scheuerman. 2020. Bringing the People Back In: Contesting Benchmark Machine Learning Datasets. <http://arxiv.org/abs/2007.07399>
- [25] Nathalie Diberardino, Clair Baleshta, and Luke Stark. 2024. Algorithmic Harms and Algorithmic Wrongs. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Rio de Janeiro Brazil, 1725–1732. <https://doi.org/10.1145/3630106.3659001>
- [26] Isak Engdahl. 2024. Agreements ‘in the wild’: Standards and alignment in machine learning benchmark dataset construction. *Big Data & Society* 11, 2 (June 2024), 20539517241242457. <https://doi.org/10.1177/20539517241242457>
- [27] Kawin Ethayarajh and Dan Jurafsky. 2021. Utility is in the Eye of the User: A Critique of NLP Leaderboards. <http://arxiv.org/abs/2009.13888>
- [28] European Union. 2022. Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and amending Directive 2000/31/EC (Digital Services Act).
- [29] European Union. 2024. First Draft of the General-Purpose AI Code of Practice published, written by independent experts.
- [30] European Union. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (Artificial Intelligence Act).
- [31] European Union. 2024. Second Draft of the General-Purpose AI Code of Practice published, written by independent experts.
- [32] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021. Datasheets for Datasets. <http://arxiv.org/abs/1803.09010>
- [33] Sebastian Gehrmann, Elizabeth Clark, and Thibault Sellam. 2023. Repairing the Cracked Foundation: A Survey of Obstacles in Evaluation Practices for Generated Text. *Journal of Artificial Intelligence Research* 77 (May 2023), 103–166. <https://doi.org/10.1613/jair.1.13715>
- [34] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. 2020. Short-cut Learning in Deep Neural Networks. *Nature Machine Intelligence* 2, 11 (Nov. 2020), 665–673. <https://doi.org/10.1038/s42256-020-00257-z>
- [35] Aryo Pradipta Gema, Joshua Ong Jun Leang, Giwon Hong, Alessio Devoto, Alberto Carlo Maria Mancino, Rohit Saxena, Xuanli He, Yu Zhao, Xiaotang Du, Mohammad Reza Ghasemi Madani, Claire Barale, Robert McHardy, Joshua Harris, Jean Kaddour, Emile van Krieken, and Pasquale Minervini. 2024. Are We Done with MMLU? <https://doi.org/10.48550/arXiv.2406.04127>
- [36] Emilia Gomez, Porcaro Lorenzo, Pedro Frau Amar, and Joao Vinagre. 2024. *Diversity in Artificial Intelligence Conferences*. Publications Office of the European Union JRC137550. Publications Office. <https://data.europa.eu/doi/10.2760/796551>
- [37] Ryan Greenblatt, Carson Denison, Benjamin Wright, Fabien Roger, Monte MacDiarmid, Sam Marks, Johannes Treutlein, Tim Belonax, Jack Chen, David Duvenaud, Akbir Khan, Julian Michael, Sören Mindermann, Ethan Perez, Linda Petrini, Jonathan Uesato, Jared Kaplan, Buck Shlegeris, Samuel R. Bowman, and Evan Hubinger. 2024. Alignment faking in large language models. <https://doi.org/10.48550/arXiv.2412.14093>
- [38] Gabriel Grill. 2024. Constructing Capabilities: The Politics of Testing Infrastructures for Generative AI. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Rio de Janeiro Brazil, 1838–1849. <https://doi.org/10.1145/3630106.3659009>
- [39] Philipp Guldemann, Alexander Spiridonov, Robin Staab, Nikola Jovanović, Mark Vero, Velko Vechev, Anna Gueorgieva, Mislav Balunović, Nikola Konstantinov, Pavol Bielik, Petar Tsankov, and Martin Vechev. 2024. COMPL-AI Framework: A Technical Interpretation and LLM Benchmarking Suite for the EU Artificial Intelligence Act. <http://arxiv.org/abs/2410.07959>
- [40] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring Massive Multitask Language Understanding. <http://arxiv.org/abs/2009.03300>
- [41] J.L. Henning. 2000. SPEC CPU2000: measuring CPU performance in the New Millennium. *Computer* 33, 7 (2000), 28–35. <https://doi.org/10.1109/2.869367>
- [42] Sara Hooker. 2024. On the Limitations of Compute Thresholds as a Governance Strategy. <http://arxiv.org/abs/2407.05694>
- [43] Ben Hutchinson, Negar Rostamzadeh, Christina Greer, Katherine Heller, and Vinodkumar Prabhakaran. 2022. Evaluation Gaps in Machine Learning Practice. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Seoul Republic of Korea, 1859–1876. <https://doi.org/10.1145/3531146.3533233>
- [44] Dietmar Jannach and Christine Bauer. 2020. Escaping the McNamara Fallacy: Toward More Impactful Recommender Systems Research. *AI Magazine* 41, 4 (Dec. 2020), 79–95. <https://doi.org/10.1609/aimag.v41i4.5312>
- [45] Florian Jatón. 2021. *The Constitution of Algorithms: Ground-Truthing, Programming, Formulating*. The MIT Press, Cambridge.

- [46] Edward B Kang. 2023. Ground truth tracings (GTT): On the epistemic limits of machine learning. *Big Data & Society* 10, 1 (Jan. 2023), 20539517221146122. <https://doi.org/10.1177/20539517221146122>
- [47] Shachar Kaufman, Saharon Rosset, Claudia Perlich, and Ori Stitelman. 2012. Leakage in data mining: Formulation, detection, and avoidance. *ACM Trans. Knowl. Discov. Data* 6, 4, Article 15 (Dec. 2012), 21 pages. <https://doi.org/10.1145/2382577.2382579>
- [48] Jon Keegan. 2024. Everyone Is Judging AI by These Tests. But Experts Say They're Close to Meaningless. *The Markup* (July 2024). <https://themarkup.org/artificial-intelligence/2024/07/17/everyone-is-judging-ai-by-these-tests-but-experts-say-theyre-close-to-meaningless>
- [49] Os Keyes and Jeanie Austin. 2022. Feeling fixes: Mess and emotion in algorithmic audits. *Big Data & Society* 9, 2 (July 2022), 2053951722113772. <https://doi.org/10.1177/2053951722113772>
- [50] Bernard Koch, Emily Denton, Alex Hanna, and Jacob G. Foster. 2021. Reduced, Reused and Recycled: The Life of a Dataset in Machine Learning Research. <http://arxiv.org/abs/2112.01716>
- [51] Bernard J. Koch and David Peterson. 2024. From Protoscience to Epistemic Monoculture: How Benchmarking Set the Stage for the Deep Learning Revolution. <http://arxiv.org/abs/2404.06647>
- [52] Travis LaCroix and Alexandra Sasha Luccioni. 2022. Metaethical Perspectives on 'Benchmarking' AI Ethics. <http://arxiv.org/abs/2204.05151>
- [53] Patrick Lewis, Pontus Stenetorp, and Sebastian Riedel. 2021. Question and Answer Test-Train Overlap in Open-Domain Question Answering Datasets. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, Paola Merlo, Jorg Tiedemann, and Reut Tsarfaty (Eds.). Association for Computational Linguistics, Online, 1000–1008. <https://doi.org/10.18653/v1/2021.eacl-main.86>
- [54] Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Cosgrove, Christopher D. Manning, Christopher Ré, Diana Acosta-Navas, Drew A. Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue Wang, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mert Yuksekgonul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. 2023. Holistic Evaluation of Language Models. <https://doi.org/10.48550/arXiv.2211.09110>
- [55] Q. Vera Liao and Ziang Xiao. 2023. Rethinking Model Evaluation as Narrowing the Socio-Technical Gap. <http://arxiv.org/abs/2306.03100>
- [56] Thomas Liao, Rohan Taori, Inioluwa Deborah Raji, and Ludwig Schmidt. 2021. Are We Learning Yet? A Meta Review of Evaluation Failures Across Machine Learning. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*. <https://openreview.net/forum?id=mPducS1MsEK>
- [57] Pengfei Liu, Jinlan Fu, Yang Xiao, Weizhe Yuan, Shuaicheng Chang, Junqi Dai, Yixin Liu, Zihuiwen Ye, Zi-Yi Dou, and Graham Neubig. 2021. ExplainaBoard: An Explainable Leaderboard for NLP. <http://arxiv.org/abs/2104.06387>
- [58] Dieuwertje Luitse, Tobias Blanke, and Thomas Poell. 2024. AI competitions as infrastructures of power in medical imaging. *Information, Communication & Society* (March 2024), 1–22. <https://doi.org/10.1080/1369118X.2024.2334393>
- [59] Inbal Magar and Roy Schwartz. 2022. Data Contamination: From Memorization to Exploitation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (Eds.). Association for Computational Linguistics, Dublin, Ireland, 157–165. <https://doi.org/10.18653/v1/2022.acl-short.18>
- [60] Nicolas Malevé. 2023. Practices of Benchmarking: Vulnerability in the Computer Vision Pipeline. *photographies* 16, 2 (May 2023), 173–189. <https://doi.org/10.1080/17540763.2023.2189159>
- [61] Noortje Marres and David Stark. 2020. Put to the test: For a new sociology of testing. *The British Journal of Sociology* 71, 3 (June 2020), 423–443. <https://doi.org/10.1111/1468-4446.12746>
- [62] Peter Mattson, Christine Cheng, Gregory Damos, Cody Coleman, Paulius Micikevicius, David Patterson, Hanlin Tang, Gu-Yeon Wei, Peter Bailis, Victor Bittorf, David Brooks, Dehao Chen, Debo Dutta, Udit Gupta, Kim Hazelwood, Andy Hock, Xinyuan Huang, Daniel Kang, David Kanter, Naveen Kumar, Jeffery Liao, Deepak Narayanan, Tayo Oguntebi, Gennady Pekhimenko, Lillian Pentecost, Vijay Janapa Reddi, Taylor Robie, Tom St John, Carole-Jean Wu, Lingjie Xu, Cliff Young, and Matei Zaharia. 2020. MLPerf Training Benchmark. In *Proceedings of Machine Learning and Systems*, I. Dhillon, D. Papailiopoulos, and V. Sze (Eds.), Vol. 2. 336–349. https://proceedings.mlsys.org/paper_files/paper/2020/file/411e39b117e885341f25efb8912945f7-Paper.pdf
- [63] Peter Mattson, Vijay Janapa Reddi, Christine Cheng, Cody Coleman, Greg Damos, David Kanter, Paulius Micikevicius, David Patterson, Guenther Schmuelling, Hanlin Tang, Gu-Yeon Wei, and Carole-Jean Wu. 2020. MLPerf: An Industry Standard Benchmark Suite for Machine Learning Performance. *IEEE Micro* 40, 2 (2020), 8–16. <https://doi.org/10.1109/MM.2020.2974843>
- [64] Timothy R. McIntosh, Teo Susnjak, Nalin Arachchilage, Tong Liu, Paul Watters, and Malka N. Halgamuge. 2024. Inadequacies of Large Language Model Benchmarks in the Era of Generative Artificial Intelligence. <http://arxiv.org/abs/2402.09880>
- [65] Julian Michael, Ari Holtzman, Alicia Parrish, Aaron Mueller, Alex Wang, Angelica Chen, Divyam Madaan, Nikita Nangia, Richard Yuanzhe Pang, Jason Phang, and Samuel R. Bowman. 2022. What Do NLP Researchers Believe? Results of the NLP Community Metasurvey. <http://arxiv.org/abs/2208.12852>
- [66] Marc Miltenberger, Steven Arzt, Philipp Holzinger, and Julius Näumann. 2023. Benchmarking the Benchmarks. In *Proceedings of the ACM Asia Conference on Computer and Communications Security*. ACM, Melbourne VIC Australia, 387–400. <https://doi.org/10.1145/3579856.3582830>
- [67] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 220–229.

- <https://doi.org/10.1145/3287560.3287596>
- [68] Moran Mizrahi, Guy Kaplan, Dan Malkin, Rotem Dror, Dafna Shahaf, and Gabriel Stanovsky. 2024. State of What Art? A Call for Multi-Prompt LLM Evaluation. <https://doi.org/10.48550/arXiv.2401.00595>
- [69] Dylan Mulvin. 2021. *Proxies: The Cultural Work of Standing In*. The MIT Press, Cambridge.
- [70] Arvind Narayanan and Sayash Kapoor. 2023. Evaluating LLMs Is a Minefield. https://www.cs.princeton.edu/~arvindn/talks/evaluating_llms_minefield/
- [71] Arvind Narayanan and Sayash Kapoor. 2023. GPT-4 and professional benchmarks: the wrong answer to the wrong question. <https://www.aisnakeoil.com/p/gpt-4-and-professional-benchmarks>
- [72] Milad Nasr, Nicholas Carlini, Jonathan Hayase, Matthew Jagielski, A. Feder Cooper, Daphne Ippolito, Christopher A. Choquette-Choo, Eric Wallace, and Katherine Lee. 2023. Extracting Training Data from ChatGPT. <https://not-just-memorization.github.io/extracting-training-data-from-chatgpt.html?ref=404media.co>
- [73] Milad Nasr, Nicholas Carlini, Jonathan Hayase, Matthew Jagielski, A. Feder Cooper, Daphne Ippolito, Christopher A. Choquette-Choo, Eric Wallace, Florian Tramèr, and Katherine Lee. 2023. Scalable Extraction of Training Data from (Production) Language Models. <http://arxiv.org/abs/2311.17035>
- [74] Luke Oakden-Rayner, Jared Dunnmon, Gustavo Carneiro, and Christopher Ré. 2019. Hidden Stratification Causes Clinically Meaningful Failures in Machine Learning for Medical Imaging. <http://arxiv.org/abs/1909.12475>
- [75] Victor Ojewale, Ryan Steed, Briana Vecchione, Abeba Birhane, and Inioluwa Deborah Raji. 2024. Towards AI Accountability Infrastructure: Gaps and Opportunities in AI Audit Tooling. <http://arxiv.org/abs/2402.17861>
- [76] Will Orr and Kate Crawford. 2024. Building Better Datasets: Seven Recommendations for Responsible Design from Dataset Creators. <http://arxiv.org/abs/2409.00252>
- [77] Will Orr and Kate Crawford. 2024. The social construction of datasets: On the practices, processes, and challenges of dataset creation for machine learning. *New Media & Society* 26, 9 (2024), 4955–4972.
- [78] Will Orr and Edward B. Kang. 2024. AI as a Sport: On the Competitive Epistemologies of Benchmarking. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Rio de Janeiro Brazil, 1875–1884. <https://doi.org/10.1145/3630106.3659012>
- [79] Simon Ott, Adriano Barbosa-Silva, Kathrin Blagec, Jan Brauner, and Matthias Samwald. 2022. Mapping global dynamics of benchmark creation and saturation in artificial intelligence. *Nature Communications* 13, 1 (Nov. 2022), 6793. <https://doi.org/10.1038/s41467-022-34591-0>
- [80] Oxford English Dictionary. 2017. *Benchmark. Meaning and Use*. https://www.oed.com/dictionary/benchmark_n?tab=meaning_and_use&tl=true
- [81] Lorenzo Pacchiardi, Marko Tesic, Lucy G. Cheke, and José Hernández-Orallo. 2024. Leaving the barn door open for Clever Hans: Simple features predict LLM benchmark answers. <http://arxiv.org/abs/2410.11672>
- [82] Jaihyun Park and Sullam Jeoung. 2022. Raison d’être of the benchmark dataset: A Survey of Current Practices of Benchmark Dataset Sharing Platforms. In *Proceedings of NLP Power! The First Workshop on Efficient Benchmarking in NLP*. Association for Computational Linguistics, Dublin, Ireland, 1–10. <https://doi.org/10.18653/v1/2022.nlppower-1.1>
- [83] Amandalynne Paullada, Inioluwa Deborah Raji, Emily M. Bender, Emily Denton, and Alex Hanna. 2021. Data and its (dis)contents: A survey of dataset development and use in machine learning research. *Patterns* 2, 11 (Nov. 2021), 100336. <https://doi.org/10.1016/j.patter.2021.100336>
- [84] Rolf Pfister and Hansueli Jud. 2025. Understanding and Benchmarking Artificial Intelligence: OpenAI’s o3 Is Not AGI. <http://arxiv.org/abs/2501.07458>
- [85] Trevor Pinch. 1993. “Testing - One, Two, Three ... Testing!”: Toward a Sociology of Testing. *Science, Technology, & Human Values* 18, 1 (Jan. 1993), 25–41. <https://doi.org/10.1177/016224399301800103>
- [86] Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson. 2023. Fine-tuning Aligned Language Models Compromises Safety, Even When Users Do Not Intend To! <http://arxiv.org/abs/2310.03693>
- [87] Inioluwa Deborah Raji, Emily Denton, Emily M. Bender, Alex Hanna, and Amandalynne Paullada. 2021. AI and the Everything in the Whole Wide World Benchmark. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*. <https://openreview.net/forum?id=j6NxpQbREA1>
- [88] Maribeth Rauh, Nahema Marchal, Arianna Manzini, Lisa Anne Hendricks, Ramona Comanescu, Canfer Akbulut, Tom Stepleton, Juan Mateos-Garcia, Stevie Bergman, Jackie Kay, Conor Griffin, Ben Bariach, Iason Gabriel, Verena Rieser, William Isaac, and Laura Weidinger. 2024. Gaps in the Safety Evaluation of Generative AI. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* 7 (Oct. 2024), 1200–1217. <https://doi.org/10.1609/aies.v7i1.31717>
- [89] Richard Ren, Steven Basart, Adam Khoja, Alice Gatti, Long Phan, Xuwang Yin, Mantas Mazeika, Alexander Pan, Gabriel Mukobi, Ryan Hwang Kim, Stephen Fitz, and Dan Hendrycks. 2024. Safetywashing: Do AI Safety Benchmarks Actually Measure Safety Progress?. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*. <https://openreview.net/forum?id=YagfTP3RK6>
- [90] Anka Reuel, Amelia Hardy, Chandler Smith, Max Lamparth, Malcolm Hardy, and Mykel Kochenderfer. 2024. BetterBench: Assessing AI Benchmarks, Uncovering Issues, and Establishing Best Practices. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*. <https://openreview.net/forum?id=hcOq2buakM>
- [91] Manley Roberts, Himanshu Thakur, Christine Herlihy, Colin White, and Samuel Dooley. 2023. Data Contamination Through the Lens of Time. <http://arxiv.org/abs/2310.10628>
- [92] Pedro Rodriguez, Joe Barrow, Alexander Miserlis Hoyle, John P. Lalor, Robin Jia, and Jordan Boyd-Graber. 2021. Evaluation Examples are not Equally Informative: How should that change NLP Leaderboards?. In *Proceedings of the 59th Annual Meeting of the Association for Computational*

- Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, Online, 4486–4503. <https://doi.org/10.18653/v1/2021.acl-long.346>
- [93] Paul Röttger, Fabio Pernisi, Bertie Vidgen, and Dirk Hovy. 2024. SafetyPrompts: a Systematic Review of Open Datasets for Evaluating and Improving Large Language Model Safety. <http://arxiv.org/abs/2404.05399>
- [94] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. 2021. “Everyone wants to do the model work, not the data work”: Data Cascades in High-Stakes AI. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–15. <https://doi.org/10.1145/3411764.3445518>
- [95] 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 (Oct. 2021), 1–37. <https://doi.org/10.1145/3476058>
- [96] David Schlangen. 2020. Targeting the Benchmark: On Methodology in Current Natural Language Processing Research. <http://arxiv.org/abs/2007.04792>
- [97] D Sculley, Jasper Snoek, Ali Rahimi, and Alex Wiltschko. 2018. Winner’s Curse? On Pace, Progress, and Empirical Rigor. Vancouver, BC, Canada. <https://openreview.net/pdf?id=rJWF0Fywf>
- [98] Andrew D. Selbst, Danah Boyd, Sorelle A. Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and Abstraction in Sociotechnical Systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM, Atlanta GA USA, 59–68. <https://doi.org/10.1145/3287560.3287598>
- [99] Shilad Sen, Margaret E. Giesel, Rebecca Gold, Benjamin Hillmann, Matt Lesicko, Samuel Naden, Jesse Russell, Zixiao (Ken) Wang, and Brent Hecht. 2015. Turkers, Scholars, “Arafat” and “Peace”: Cultural Communities and Algorithmic Gold Standards. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, Vancouver BC Canada, 826–838. <https://doi.org/10.1145/2675133.2675285>
- [100] Jan Simson, Alessandro Fabris, and Christoph Kern. 2024. Lazy Data Practices Harm Fairness Research. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Rio de Janeiro Brazil, 642–659. <https://doi.org/10.1145/3630106.3658931>
- [101] Jessie J. Smith, Saleema Amershi, Solon Barocas, Hanna Wallach, and Jennifer Wortman Vaughan. 2022. REAL ML: Recognizing, Exploring, and Articulating Limitations of Machine Learning Research. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. 587–597. <https://doi.org/10.1145/3531146.3533122>
- [102] Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shueb, Abubakar Abid, Adam Fisch, Adam R. Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, Agnieszka Kluska, Aitor Lewkowycz, Akshat Agarwal, Alethea Power, Alex Ray, Alex Warstadt, et al. 2023. Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models. <http://arxiv.org/abs/2206.04615>
- [103] Isabelle Stengers. 2018. *Another science is possible: a manifesto for slow science*. Polity press, Cambridge.
- [104] Matthew Stewart. 2023. The Olympics of AI: Benchmarking Machine Learning Systems. <https://towardsdatascience.com/the-olympics-of-ai-benchmarking-machine-learning-systems->
- [105] Marilyn Strathern. 1997. ‘Improving ratings’: audit in the British University system. *European Review* 5, 3 (July 1997), 305–321. [https://doi.org/10.1002/\(SICI\)1234-981X\(199707\)5:3<305::AID-EURO184>3.0.CO;2-4](https://doi.org/10.1002/(SICI)1234-981X(199707)5:3<305::AID-EURO184>3.0.CO;2-4)
- [106] Arjun Subramonian, Xingdi Yuan, Hal Daumé III, and Su Lin Blodgett. 2023. It Takes Two to Tango: Navigating Conceptualizations of NLP Tasks and Measurements of Performance. <http://arxiv.org/abs/2305.09022>
- [107] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. BEIR: A Heterogeneous Benchmark for Zero-shot Evaluation of Information Retrieval Models. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*. <https://openreview.net/forum?id=wCu6T5xFje>
- [108] The White House. 2023. Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.
- [109] Nanna Bonde Thylstrup, Kristian Bondo Hansen, Mikkel Flyverbom, and Louise Amoore. 2022. Politics of data reuse in machine learning systems: Theorizing reuse entanglements. *Big Data & Society* 9, 2 (July 2022), 20539517221139785. <https://doi.org/10.1177/20539517221139785>
- [110] Kushal Tirumala, Aram Markosyan, Luke Zettlemoyer, and Armen Aghajanyan. 2022. Memorization Without Overfitting: Analyzing the Training Dynamics of Large Language Models. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 38274–38290. https://proceedings.neurips.cc/paper_files/paper/2022/file/fa0509f4dab6807e2cb465715bf2d249-Paper-Conference.pdf
- [111] Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Andrew Ilyas, and Aleksander Madry. 2020. From ImageNet to Image Classification: Contextualizing Progress on Benchmarks. <http://arxiv.org/abs/2005.11295>
- [112] UK Parliament. 2023. Online Safety Act 2023.
- [113] US Department of Commerce. 2025. Framework for Artificial Intelligence Diffusion.
- [114] Keyon Vafa, Justin Y. Chen, Ashesh Rambachan, Jon Kleinberg, and Sendhil Mullainathan. 2024. Evaluating the World Model Implicit in a Generative Model. <http://arxiv.org/abs/2406.03689>
- [115] Laura Weidinger, Maribeth Rauh, Nahema Marchal, Arianna Manzini, Lisa Anne Hendricks, Juan Mateos-Garcia, Stevie Bergman, Jackie Kay, Conor Griffin, Ben Bariach, Iason Gabriel, Verena Rieser, and William Isaac. 2023. Sociotechnical Safety Evaluation of Generative AI Systems. <http://arxiv.org/abs/2310.11986>
- [116] Teun van der Weij, Felix Hofstätter, Ollie Jaffe, Samuel F. Brown, and Francis Rhys Ward. 2024. AI Sandbagging: Language Models can Strategically Underperform on Evaluations. <http://arxiv.org/abs/2406.07358>
- [117] Cheng Xu, Shuhao Guan, Derek Greene, and M.-Tahar Kechadi. 2024. Benchmark Data Contamination of Large Language Models: A Survey. <http://arxiv.org/abs/2406.04244>

- [118] Lifan Yuan, Yangyi Chen, Ganqu Cui, Hongcheng Gao, Fangyuan Zou, Xingyi Cheng, Heng Ji, Zhiyuan Liu, and Maosong Sun. 2023. Revisiting Out-of-distribution Robustness in NLP: Benchmark, Analysis, and LLMs Evaluations. *37th Conference on Neural Information Processing Systems (NeurIPS 2023) Track on Datasets and Benchmarks (2023)*.
- [119] Andy K. Zhang, Kevin Klyman, Yifan Mai, Yoav Levine, Yian Zhang, Rishi Bommasani, and Percy Liang. 2024. Language model developers should report train-test overlap. <http://arxiv.org/abs/2410.08385>
- [120] Lin Zhijia. 2024. Top LLMs in China and the U.S. Only 5 Months Apart: Kai-Fu Lee. <https://en.tmtpost.com/post/7289212>

This figure "acm-jdslogo.png" is available in "png" format from:

<http://arxiv.org/ps/2502.06559v1>

This figure "image.png" is available in "png" format from:

<http://arxiv.org/ps/2502.06559v1>