

Scaling Laws Do Not Scale

Fernando Diaz^{1,2}, Michael Madaio¹

¹Google Research

²Carnegie Mellon University
diazf@acm.org, madaiom@google.com

Abstract

Recent work has advocated for training AI models on ever-larger datasets, arguing that as the size of a dataset increases, the performance of a model trained on that dataset will correspondingly increase (referred to as “scaling laws”). In this paper, we draw on literature from the social sciences and machine learning to critically interrogate these claims. We argue that this scaling law relationship depends on metrics used to measure performance that may not correspond with how different groups of people perceive the quality of models’ output. As the size of datasets used to train large AI models grows and AI systems impact ever larger groups of people, the number of distinct communities represented in training or evaluation datasets grows. It is thus even more likely that communities represented in datasets may have values or preferences not reflected in (or at odds with) the metrics used to evaluate model performance in scaling laws. Different communities may also have values in tension with each other, leading to difficult, potentially irreconcilable choices about metrics used for model evaluations—threatening the validity of claims that model performance is improving at scale. We end the paper with implications for AI development: that the motivation for scraping ever-larger datasets may be based on fundamentally flawed assumptions about model performance. That is, models may not, in fact, continue to improve as the datasets get larger—at least not for all people or communities impacted by those models. We suggest opportunities for the field to rethink norms and values in AI development, resisting claims for universality of large models, fostering more local, small-scale designs, and other ways to resist the impetus towards scale in AI.

Introduction

In the context of multilingual language modeling, the *curse of multilinguality* refers to the degradation of model performance across languages as the diversity and size of the training data grows (Conneau et al. 2020; Chang et al. 2023). These results suggest that bounded model capacity forces competition amongst languages for model parameters, leading to degradation of performance for *all* languages. In situations where languages are not uniformly represented, low resource languages tend to suffer more (Wu and Dredze

2020). Similar results have also been observed in the context of mixed-modal models (Aghajanyan et al. 2023). We argue that the curse of multilinguality phenomena extends beyond language and modalities to include the values of different communities or sub-groups of people. For large pre-trained models, model developers and evaluators often do not know about the diverse space of sub-groups, nor how to measure a model’s performance for each community, making any claims of improved model performance limited to the values of an over-represented, high-resource sub-group.

We ground our analysis in the development of *scaling laws* for AI: the relationship between machine learning model performance (as measured by some evaluation metric) and model design variables (e.g., dataset size, number of model parameters, compute) (Kaplan et al. 2020). Existing research suggests that this relationship is superlinear for deep learning models, and it has been used to justify the collection of ever-larger datasets used to train large language models (Rae et al. 2021; Chowdhery et al. 2022).

However, the validity of these scaling laws relies on the extent to which the metrics used in scaling law analyses actually reflect the quality of a given model’s performance for all groups of people impacted by large AI models deployed at scale. In the context of AI systems that are used by or impact people, an evaluation metric (the dependent variable underlying many scaling laws) is intended to measure the performance or quality of the system *for those people*. The quantity computed by an evaluation metric is shaped both by its underlying mathematical assumptions and by the sampling procedures used to collect data to compute it. On one hand, the *mathematical form* of an evaluation metric ideally reflects what a system designer thinks is important to users or others impacted by the system, encoding assumptions about users’ beliefs, values, and behavior. On the other hand, the *evaluation data* used to compute the metric is meant to capture some sense of the quality of that model by sampling a particular population of interest. In combination, the mathematical form of a metric and the associated data used to compute it estimate the performance of a model for a specific population, as is especially salient in emerging learned evaluation metrics (Rei et al. 2020; Stanojević and Sima’an 2014).

Therefore, when designing for a sufficiently diverse population of people, measuring performance with a single, uni-

versal metric may lead to misplaced trust in models' quality. Populations are composed of *subpopulations*: members of particular identity groups, including demographic groups (e.g., gender, racial or ethnic groups, religion, caste, national identity, disability status), as well as other cultural or sociopolitical groups. We argue that as the size of the population being designed for increases, the number of subpopulations present in any evaluation is likely to increase. Samples used to evaluate model performance will capture some set of communities in the social context in which the data was collected.

Increasing the number of impacted people and the evaluation dataset size may thus increase the number of subpopulations present in the data and correspondingly increase the range of values that should be considered, potentially leading to values tensions (Miller et al. 2007; Ghassemi and Mohamed 2022; Birhane et al. 2022; Durmus et al. 2023; Sorensen et al. 2023; Varshney 2023). Beyond the well-known challenges of developing AI systems to be used by diverse subpopulations—such as evaluating algorithmic unfairness in model performance with respect to a fixed evaluation metric (Hardt, Price, and Srebro 2016)—we argue that diverse subpopulations are likely to have different and potentially conflicting notions of 'good' (Green 2019) as instantiated in both the mathematical form and data behind an evaluation metric, which machine learning evaluations using single metrics often assume to be universal. While algorithmic unfairness surfaces the systematic variation of the true relationship between features and a single universal metric across subpopulations, we contend that the metrics themselves may vary across subpopulations.

The current lack of attention to the subpopulations and communities represented in scaling law evaluation datasets poses major challenges for proponents of AI scaling laws—and presents potential risks for all those impacted by the deployment of large models. Despite claims that a larger training dataset will lead to improved model performance, when such models are deployed at scale, the larger numbers of people who are impacted by such systems or included in the evaluation dataset may lead to breakdowns in model performance for different communities (Bender et al. 2021). Different communities of users or impacted stakeholders may necessitate evaluating with different metrics (due to different behaviors or values), each of which may be in conflict with each other, or in conflict with commonly used evaluation metrics.

We draw on empirical results from machine learning and scholarship from the social sciences to interrogate the validity of claims made about scaling laws for training dataset size and model performance when considering their use in models deployed at scale, on global populations whose values may not be reflected in current performance metrics, or which may be in irreconcilable tension with values of other communities. These observations suggest that scaling laws, in the (perhaps misguided) pursuit of general principles, can obscure systematic under-performance for subpopulations and communities of people impacted by these models.

We thus propose that proponents of scaling laws reconsider their claims to universality of model performance

across increasingly global populations of users of such large models. This may involve rethinking choices about evaluation metrics and the culturally-situated nature of evaluation datasets used to compute them; tempering claims for general applicability of scaling laws' evaluations of model performance; and making intentional choices about which communities are involved in evaluating models and how their values may be reflected in evaluation choices, including investigating how to resolve tensions between multiple communities' values when evaluating large models deployed at global scales. This paper contributes to a growing body of scholarship that explores the impacts of large models on global populations.

Scaling Laws

In order to scrutinize scaling laws from the perspective of evaluation metrics, we first review the relevant concepts from AI evaluation and scaling law literature. We will introduce some concepts from measurement modeling (Hand 2010; Jacobs and Wallach 2021) and introduce some notation that we will refer to throughout our analysis of scaling laws.

Performance Metrics Are Imperfect Proxies for Model Quality

Evaluation of AI systems involves computing and comparing quantitative measures of performance of a system on a task. In offline settings, including laboratory or benchmark experiments, researchers use evaluation metrics based on data labeled through dedicated annotators. In online settings, including deployed AI systems in production environments, organizations use evaluation metrics based on logged behavior data. In both settings, evaluation metrics play a critical role in guiding high-level research and model development decisions as well as more granular parameter optimization (Cohen 1995; Kohavi, Tang, and Xu 2020; Joachims 2002; Grotov and de Rijke 2016). In particular, offline metrics are often used—implicitly or not—as proxies or predictors of online metrics (Zheng 2015; Suresh and Gutttag 2021; Rudin and Wagstaff 2014). In line with modern AI paradigms, we focus on the evaluation of models trained on a set of data. We represent a model trained on a dataset \mathcal{D} as $\pi(\mathcal{D})$. For the purpose of our analysis, we are interested in the relationship between data and model performance and therefore will not specify the input or output space of $\pi(\mathcal{D})$, nor do we care about the specific functional form or design of $\pi(\mathcal{D})$. We use Π to refer to the space of all trained models, of which $\pi(\mathcal{D})$ is one member. For clarity, we will sometimes refer to a model as π , even though it is always the outcome of a training procedure and data.

Generic metrics like accuracy or squared error are often adopted in machine learning contexts to evaluate the quality of the model output. These metrics are simple, well-understood, and usually amenable to direct optimization. For many AI systems, however, generic metrics do not capture how *people* use the system output. For instance, for a language model used for question-answering, how reflective is the measure of 'accuracy of predicting the next word' of the

underlying *user goal* of ‘understanding the answer to a question’? For a recommender system, how reflective is the measure of ‘accuracy of predicting a rating’ of the underlying *user goal* of ‘discovery of new, relevant content’? There is subtlety in how the system output is used by people, which is lost when performance is measured with generic metrics. This is why, in applied contexts, domain-specific metrics are usually developed. Indeed, a variety of areas of research, including natural language processing (e.g., BLEU (Reiter 2018; Post 2018), ROUGE (Ganesan 2018)), search and recommendation (Chandar, Diaz, and St. Thomas 2020a), and more (Raji et al. 2021), have adopted families of metrics informed by technology use (for a critique of NLP metrics, see Subramonian et al. 2023).

Even within a given domain, the quality of a model’s output may be tightly coupled with a user’s specific task; for example, the quality of a predictive typing application may be related to how useful users find it in effectively completing a writing task—but this quality may differ greatly between different use cases for the same task, such as (for instance) informal messages and creative writing tasks compared to professional communication or scientific writing tasks. For AI systems, we are thus interested in evaluation metrics that measure—explicitly or not—the quality of a system’s output (e.g., predictions, decisions, recommendations) for a given *population of people* for a given *task*.

Drawing on the measurement modeling literature (Jacobs and Wallach 2021; Hand 2010), we refer to the unobservable domain- and task-specific notion of model output quality as a *construct*, or, the true performance of a system when used by a specific population for a specific task. While in many cases, a construct is *completely* unobservable, in others it may merely be very expensive to collect or more accurately estimate (e.g., downstream revenue effects; user satisfaction). We define $\mu^*(\mathcal{U}, \pi(\mathcal{D}))$ as the *latent* scalar value associated with the quality of $\pi(\mathcal{D})$ for a population \mathcal{U} . This population could be compact and well-defined (e.g., ‘coworkers in a specific academic department’) or vague and general (e.g., ‘any person with access to a computer, regardless of location, race, ethnicity, age, religion, gender identity, sexual orientation, disability, or economic status’). Although it may be difficult or impossible to observe the latent construct—both because μ^* and \mathcal{U} are unobserved—we can “operationalize” it using an observable *evaluation metric* as a proxy (Jacobs and Wallach 2021; Subramonian et al. 2023). For example, we can evaluate the quality of a text prediction model’s output in terms of its usefulness for a given writing task by, for instance, measuring composition time, number of edits, the number of accepted text prediction suggestions, or using delayed user feedback on the quality of the text predictions (Robertson et al. 2021)—but all of these are only ever proxies (some better than others) for the latent construct of interest, μ^* .

To make this clear, let $\mu(U, \pi(\mathcal{D}))$ be the evaluation metric for a sample of user-associated data $U \sim \mathcal{U}$ (or, a sample U from a population \mathcal{U}). User-associated data, often referred to as an ‘evaluation dataset,’ may consist of individual data (i.e., each data point or instance is a person) or derived data (e.g., each data point is a text document written by a person).

So, an evaluation metric approximates both the functional form of μ^* as well as the associated target population \mathcal{U} . In the next section, we discuss the ways that populations might impact the validity of evaluation metrics. Let \mathcal{M} be the set of *all* evaluation metrics, (including better and worse proxies for the latent construct), that can be used for measuring system performance. From measurement modeling, the validity of an evaluation metric μ is based on the extent to which it captures the salient aspects of the latent construct μ^* it is designed to operationalize (Jacobs and Wallach 2021). The better the metric captures the construct, the better able that it is to assess the performance of a given model. Thus, the relationship between what we may want to measure about AI systems (e.g., $\mu^*(\mathcal{U}, \pi(\mathcal{D}))$) and the way we measure that (e.g., $\mu(U, \pi(\mathcal{D}))$) may not be straightforward.

Scaling Laws: Learning Curves and Power Laws

Scaling law analyses use “learning curves” to represent the relationship between system performance and training dataset size (i.e., the horizontal axis is a sequence of training datasets ordered by increasing size $|\mathcal{D}_i| < |\mathcal{D}_{i+1}|$). In general, proponents of AI scaling laws argue that $\mu^*(\mathcal{U}, \pi(\mathcal{D}_i)) < \mu^*(\mathcal{U}, \pi(\mathcal{D}_{i+1}))$ —i.e., that models trained on larger training datasets will perform better according to a given performance metric. However, such claims are often made about the unobservable construct of a model’s performance (μ^*), but they are *based on* observations from proxy metrics (μ) instead of the construct (μ^*) and population samples (\mathcal{U}) instead of the true target distribution (\mathcal{U}), as constructs and the full population may be difficult or impossible to directly observe (Ganguli et al. 2022a; Rish 2023).

The precise relationship between $|\mathcal{D}_i|$ and $\mu(U, \pi(\mathcal{D}_i))$ often follows a power law; i.e., as long as it is not bottlenecked by its capacity or access to compute resources, a model’s performance improves superlinearly as a function of training dataset size (Rosenfeld et al. 2020; Kaplan et al. 2020). Although initial results demonstrated scaling laws for natural language processing tasks, similar laws have been developed for multimodal and reinforcement learning models (Cherti et al. 2022; Hilton, Tang, and Schulman 2023). The evidence of scaling laws has motivated their usage in informing model design. They have been used as motivation for scraping ever-larger training datasets (Dodge et al. 2021), to developing quantization methods (Dettmers and Zettlemoyer 2022), to extrapolating performance from smaller datasets (Ivgi, Carmon, and Berant 2022), and to determining data minimization policies (Shanmugam et al. 2022). Despite their popularity, however, insufficient attention to the precarious relationship between latent constructs of model quality and their operationalization in performance metrics *for particular communities included in evaluation datasets* (or who are otherwise impacted by a model) poses serious questions to the validity of relying on scaling laws when evaluating deployed models at scale.

The Precarity of Metrics

In order to understand precisely how scaling laws might be compromised, we first need to review the various ways in

which metrics are far from the ‘ground truth’ they are often considered to reflect. Through a discussion and review of existing observations in the computer science and social science literature, we demonstrate that evaluation metrics are inherently contestable and precarious and that, at any one point, there may be multiple metrics and constructs in tension. We draw on and extend recent work in the responsible AI and computational social science communities that similarly recognizes that, although presented as a reliable proxy for the construct, evaluation metrics are often tenuous (Friedler, Scheidegger, and Venkatasubramanian 2016; Jacobs and Wallach 2021; Thomas and Uminsky 2022; Wagner et al. 2021).

Metric (In)compatibility

As previously discussed, an evaluation metric can be seen as an approximation of the construct of interest. A generic metric like accuracy is often adopted because it can be used without modification. But, for any given task, a construct is more complicated than suggested by a generic metric. It might be impacted by a user’s expectations, interface constraints, or normative social values, none of which are captured by a generic metric like accuracy. Because metrics are errorful proxies for constructs, any two metrics, even if they largely agree with the construct, may disagree with each other. Moreover, because constructs and their operationalizations may be contested, two metrics may disagree because each is modeling a different (though related) underlying construct (e.g., two fairness metrics may be operationalizing different understandings of fairness (Chouldechova 2017; Jacobs and Wallach 2021)). Thus, adopting a single metric for evaluating a model—including for scaling law analyses—requires validating it with respect to the construct of interest, and how that relationship may be more or less stable for various use cases, social contexts, or even different ranges of the metric’s value. Moreover, understanding the relationship *between* multiple metrics for different ranges of the metrics’ values helps designers avoid situations where a metric becomes unreliable as model performance improves.

Metric Nonstationarity

The ability of a metric to capture a construct of interest may also change over time. First, due to the complexity of sociotechnical systems made up of deployed AI systems used by or impacting people, metric development—like model development—is iterative, based on our understanding of how people and social groups engage with the technology (Chandar, Diaz, and St. Thomas 2020b), as has been found in web search evaluation (Carterette, Kanoulas, and Yilmaz 2012) and multidocument summarization evaluation (Shapira et al. 2021). At any point in time, a specific metric captures the researcher’s or practitioner’s best understanding of how to model the construct. But, because our understanding of users and the world is changing, this metric will also change over time (cf. Subramonian et al. 2023). Second, as suggested by the possible heteroscedastic relationship between metrics and constructs, the ‘appropriate metric’ at any point in time may depend on the set of systems being compared. Because model performance will

change—ideally improve—over time, metrics will become stale and require replacement with other, more appropriate metrics (Voorhees, Craswell, and Lin 2022), due to the fallibility of proxy metrics with respect to a target metric or construct (Gao, Schulman, and Hilton 2022; Skalse et al. 2022). Third, model development often exists in a broader sociotechnical context, which itself might change over time. Consider a metric such as ‘consumption time,’ used in a number of media platforms to model relevance by measuring how long someone spends consuming media. The higher the consumption time of a particular piece of media, the more relevant the item is to that user. However, improvements to bandwidth and rendering times on devices can result in substantive changes in the relationship between the consumption time and relevance. Or, acute shocks to the environment (e.g., unexpected exogenous events, such as disasters, or regulatory policy changes) can change how people, technology, and their measured values behave and, as a result, relate to a construct of interest (cf. Selbst et al. 2019). In some cases, the model itself may lead to behavioral or environmental changes that impact the usefulness of the metric. For example, users of a system may, over time, adapt their behavior in ways that compromise assumptions in the metric (e.g., how people express queries on mobile search engines (Kamvar and Baluja 2006, 2007)). Or, so-called adversaries may attempt to manipulate the system by gaming the metric (e.g., clickbait (Chen, Conroy, and Rubin 2015)).

Metric Staging

Multiple metrics may arise for pragmatic reasons. In model development, different metrics are used at different stages due to the costs involved in metric computation (Zheng 2015; Suresh and Gutttag 2021). In production systems, behavioral metrics based on real interactions with end users (e.g., user engagement metrics) may be more accurate reflections of a latent construct (e.g., due to face validity (Jacobs and Wallach 2021)), but more expensive because—in addition to potential costs of deployment or data collection—they incur a risk of harming the users, losing their trust, and having them abandon the technology altogether (Wang et al. 2023). On the other hand, in offline evaluation, metrics based on data labeled by annotators (e.g., precision, recall, simulations) are less accurate reflections of the construct than behavioral metrics. That said, they are extremely fast to compute, often reusable, and can be used in an isolated environment where unexpected or harmful outputs may be insulated from impacting users.¹ Unfortunately, offline evaluation requires an annotation budget and the development of guidelines for annotators, both of which can be prohibitive. In practice, model development involves coordinating multiple metrics: using offline evaluation and associated metrics for a large set of models before selecting a subset for production evaluation using a different set of metrics.

¹However, harmful outputs are unfortunately often not prevented from impacting annotators (cf. Gray and Suri 2019).

Metric Variation Across Subtasks

Although an AI system intended for a specific task or domain is often evaluated using a single construct or metric, in reality, there are often a diversity of specialized subtasks or subdomains. For example, both question answering systems and search engines support a wide variety of possible user information needs (Broder 2002; Murdock and Croft 2002). In turn, each subtask or subdomain has a unique construct. This may manifest from some property of the labels (e.g., contextualizing labeled data on the subtask or subdomain) or the formula for the metric itself (e.g., in the search context, reciprocal rank for navigational intents versus recall for literature review). In some situations, subtask metrics may be compatible or ‘compatible enough’, allowing for shared information when evaluating a model (Zhou et al. 2021). In other situations, two subtasks may be quite different constructs in tension (Aribandi et al. 2022). While a single, universal metric may be one way to resolve inconsistency amongst subtask or subdomain constructs or metrics, in a reinforcement learning context, Skalse and Abate (2022) demonstrate that jointly optimizing a single policy by reducing multiple rewards (i.e., metrics) to a single number is only possible in a limited set of situations. However, despite this variation across tasks, the dominant paradigm used in scaling law analyses is to optimize and evaluate for a single performance metric.

Metric Power

When metrics are used to guide research in academic communities, make decisions in industry, and evaluate performance of machine learning systems, they become social objects amongst researchers, funding agencies, engineers, product managers, and other individuals involved in the processes of machine learning research, development, deployment, and use. As such, metrics used to evaluate AI systems are always subject to (and in turn, shape) the social dynamics of the sociocultural systems in which they are embedded. When metrics are used to measure phenomena in social contexts, it is well-established that they not only enable one to understand a particular phenomena, but they also have the power (cf. Beer 2016) to change social actors’ behaviors. In other words, metrics do not simply reflect the world, but they shape it, and are shaped by it (cf. Pinch and Bijker 1984). This phenomena, known as Goodhart’s Law (Goodhart 1984; Campbell 1979; Strathern 1997), has been identified across numerous fields, including education (Strathern 1997; Griesemer 2020), economics (Mügge 2022), and organizational studies (Gray, Micheli, and Pavlov 2015). For instance, when schools use students’ test scores as a metric to evaluate teacher quality, some teachers and administrators have responded by teaching to the test or in the worst case, altering students’ test scores (Gabriel 2010). In the workplace more generally, companies have long adopted performance metrics of various sorts to measure their employees’ productivity (Ranganathan and Benson 2020) and attempt to incentivize them to be more productive (Zelizer 1996)—metrics that are increasingly based on fine-grained data from those employees, and which may shape employees’ behaviors in other ways (e.g., employees using applications

to move their cursor to simulate productivity) (Bernstein 2017). However, the specific ways that metrics impact people’s behavior are themselves shaped by the norms, culture, and organizational dynamics of the context in which they are used. For instance, Christin (2020) found that two newsrooms in France and the United States that had access to similar web traffic data about their news stories differed greatly in how the metrics they derived from those data shaped their approach to journalistic decision-making, in ways impacted by their respective organizational, professional, and cultural contexts. In addition to the culturally specific ways that metrics may impact particular social worlds, metrics may take on a social life of their own. Within research communities, metrics can have a stickiness when entire research programs develop around them (Cohen and Howe 1988).

All of this suggests that the performance metrics used to evaluate AI systems may be inherently unstable or precarious in ways that raise serious questions for the validity and robustness of scaling laws for AI, which rely on metrics that are often divorced from any particular social context.

Values Pluralism Threatens Scaling Laws at Scale

We now turn to critically interrogate scaling laws in light of the multiplicity of human values across the numerous social groups and communities impacted by models and represented in evaluation datasets. While there are existing initiatives to identify ‘inverse scaling laws’ for particular tasks (Wei, Tay, and Le 2022), our claim is that, when we consider training and evaluating with human data, scaling laws, as currently posed, are, at best, incomplete, and may be fundamentally flawed. Our claim is divided into four parts. First, that evaluation metrics reflect the composition of the evaluation dataset, which is shaped by the sampling approach used to collect that data; second, that the number of sub-groups within a given evaluation dataset grows with data size; third, those sub-groups can have incompatible values and preferences for appropriate evaluation metrics; and fourth, that the risk of that metric incompatibility grows with dataset size.

Sampling Approaches Shape Who Is Included in Evaluation Datasets

As we mentioned earlier, metric-based evaluation estimates the performance of a model when used by an *intended* population of users. Consistent with standard assumptions in machine learning, ideally both training and evaluation data are drawn from the same population using the same sampling distribution (although in practice this may not be the case). If \mathcal{U} is the population, then θ defines the sampling distribution from which we draw both the training set \mathcal{D} and the evaluation set \mathcal{U} . When describing an evaluation metric, we can describe it in terms of the sample size $|\mathcal{U}|$. For instance, benchmarking or offline experiments have a fixed sample (and sample size) determined by the collected, static evaluation data, where training and testing sets are often defined by a percentage of instances (e.g., 80% training; 20% testing). In deployed systems, the sample size varies with the number of users (and the training data size $|\mathcal{D}|$).

While we have talked about models evaluated when used by or impacting people for specific tasks, we have avoided the question of *which people* an AI system is designed for, used by, or impacts. Given a current set of users (e.g., in a deployed system), we can answer this question narrowly by evaluating a model with respect to the existing set of users of a system or application that a given model is embedded in; in this case, we say that the sampling distribution has “support” constrained to the subpopulations reflected by the current set of users. Or, we can answer this aspirationally by evaluating with respect to some future population of users; in this case, we assume that the sampling distribution has support that bounds all subpopulations present in the complete population. For the purpose of our analysis, because of the aspirational nature of claims made about AI scaling laws (Kaplan et al. 2020), we assume that the support of the sampling distribution is complete. This implies that every possible user has a nonzero probability of occurring in the training or evaluation data.

However, regardless of the sampling distribution, in any specific sample, we rarely (if ever) have reliable data for every possible user in every context. Nevertheless, the desire for such a dataset has led researchers to seek out larger sources of data on human populations (Hargittai 2015; Yoo 2015; Kitchin 2014; Anderson 2008). Debates about whether and how different groups of people may be represented by or within large datasets have persisted for decades, across multiple fields (Chasalow and Levy 2021; Bergman et al. 2023). Social scientists conducting public opinion surveys have long wrestled with what it means for their samples to be representative (and of whom they might be representative), including, for instance, examining representativeness of survey respondents from various demographics across geographic scales (e.g., cities, states, or national populations) (de Marchi and Page 2016). Many theories and methods in the social sciences grappled with the fundamental heterogeneity of human populations, as groups of people may vary based on gender, age, race/ethnicity, disability status, as well as behavior, physiology, language, religion, culture, and numerous other dimensions (Feczko et al. 2019). These dimensions of difference may themselves be more or less stable, as people may claim membership in numerous communities at different points in their lives, as such identities become more or less salient (e.g., a person leaving or joining a religion, moving to a new city or country, joining or leaving the military, and numerous other examples). In other cases, the community itself may spring into existence in response to a particular political issue (e.g., DREAMers, NIMBYs, climate change or election deniers, and more) (DiSalvo 2009).

Depending on one’s question of interest, different approaches may be needed to understand how a sample captures variation across a larger population (Feczko et al. 2019). That is, the means we use to capture these different dimensions or communities are often designed based on the goal: political surveys may be designed to capture population-level variation across politically salient demographic categories, such as gender, race/ethnicity, education level, and income, while marketing research surveys may be designed to specifically target consumer-

relevant demographics and behaviors, such as family size, technology use, spending behaviors, and more (de Marchi and Page 2016). Although many datasets about people, such as opinion polls, are intentionally designed to answer specific research questions, the advent of large-scale data collection enabled by data traces on digital platforms has led computational social scientists and others to explore how such datasets may enable them to better understand people (Yoo 2015; Kitchin 2014). Despite claims for massive datasets (in the form of ‘big data’) to usher in “the end of theory” (Anderson 2008) via datasets where, allegedly, “*n = all*” (Kitchin 2014), subsequent research has demonstrated that large-scale platform datasets are always reflections of behaviors of particular groups of people (rather than being somehow inclusive of everyone) (boyd and Crawford 2012).

As large datasets from social media platforms (e.g., Twitter, Facebook, Reddit, etc) are used to investigate research questions not only about platform use, but about social dynamics more generally, these platforms may suffer from a ‘model organism’ problem (Tufekci 2014), where claims are made about the world in general, based on data from one specific organism, such as mice or fruit flies, or one specific social media platform or dataset. In reality, there is a non-random selection of users into social media platforms (Hargittai 2015); for example, Twitter is used by less than 20% of the US population (Mitchell and Hitlin 2013). An empirical analysis of who is left out of so-called “big data” from social media found that social media users tend to be more educated, higher-income, and more technologically savvy than non-users, with substantial differences in gender and race/ethnicity across different platforms. (Hargittai 2020). Similarly, massive datasets used to train large models, such as the Colossal Clean Crawled Corpus (C4), trained on a crawl of the web, or others such as LAION (Luccioni and Viviano 2021; Birhane, Prabhu, and Kahembwe 2021), are not representative of everyone on the planet, but contain particular snapshots of particular populations (and not others) (Bender et al. 2021). For instance, some communities may not be included at all in the C4 corpus or other large web-based corpora: low-literacy communities or those who rely on radio for information and communication (Khan et al. 2017); communities with low or no technology use—or those whose technology use is primarily on mobile devices, and does not involve producing web-based content legible to web crawlers (Okeleke and Suardi 2019; Aker and Mbiti 2010); or communities who speak low-resource languages (Magueresse, Carles, and Heetderks 2020), among many others (Bender et al. 2021).

As such, it is an open question as to who, precisely, is represented in massive datasets used to train and evaluate models for scaling laws (Bergman et al. 2023; Bender et al. 2021). The composition of a dataset depends on the sampling approach used to create it—whether that is a random population sample, a sample stratified by some dimension(s) (e.g., randomized within different state populations), or a convenience or platform sample (i.e., a sample shaped by the nature of the platform used to collect the data, such as Twitter, Reddit, or the web). For all of these approaches, we cannot say with certainty precisely how many communities or

sub-groups are reflected in a given dataset without access to information about its sampling or data collection methodology. Even then, because the definition of a given community and an individual’s membership in it may be fluid over time or potentially overlapping or intersectional (Crenshaw 1990; Collins 2019; Ovalle et al. 2023b), the number of communities represented in a dataset may depend on the research questions and methods used to investigate that question.

The Number of Sub-groups in Evaluation Datasets Grows With Dataset Size

Since evaluation metrics are influenced by the sampled population U , which itself may be non-representative for a variety of reasons, often related to the sampling approach, we now discuss how, as we increase the sample size, the number of sub-groups present will also increase. Consider a theoretical sampling distribution that aims to be representative of the population \mathcal{U} . In this case, the sequence of samples U_0, U_1, \dots will, although initially missing out on many sub-groups, converge toward covering a broad set of people, ostensibly representing all sub-groups in \mathcal{U} , which is consistent with the aspirational, though unrealistic, claims for large models to benefit “all of humanity” or “the human race” (Altman 2023). The growth in heterogeneity as a function of population has similarly been studied in political science contexts (Dahl and Tufte 1973). As mentioned, the sampling frames typically used to evaluate AI systems are *not* intentionally collected to be representative of any particular sub-group or community—nor is it clear what it would even mean for a sampling frame to be representative of all sub-groups of people (Chasalow and Levy 2021).

This means that although smaller samples are likely to contain fewer sub-groups and a large enough sample may theoretically converge toward a broad, global set of sub-groups (with the caveat that some groups may not be known *a priori*, or may only emerge due to exogenous causes after data has been collected), the rate at which the number of sub-groups are encountered is dependent on properties of the sampling distributions, and the broader sampling approach taken. Given the tendency for data collection practices to be biased toward communities that are easier for the researchers collecting the data to access (Madaio et al. 2022), or datasets that are easier to scrape (Dodge et al. 2021), the rate of observing new sub-groups will be slower than a representative sample. While stratified sampling by subpopulation may suggest a possible solution, there are several reasons to be suspicious of this approach. Enumerating all possible relevant subpopulations present in the broad target population imagined for AI is a daunting task, in part because of the large number of subpopulations but also because of the fact that they can emerge and dissolve over time.

More realistically, if sampling is done by ‘organic user growth,’ as is typical in production settings, the sampling distribution itself is changing as sample size increases. Consider deployed systems, where early adopters of the technology will not be representative of later users. For example, user growth on social media platforms tends to occur non-uniformly within and across national boundaries (Poushter, Bishop, and Chwe 2018; Perrin 2015; Wilkin-

son and Thelwall 2010). Assuming consistent adoption and growth, we can use a model of nested subpopulation support, $\text{supp}(\theta_0) \subseteq \text{supp}(\theta_1) \subseteq \dots \subseteq \mathcal{U}$. As suggested by social media adoption, the growth in support is structured: some populations tend to adopt before others due to homophily.

These observations suggest that, regardless of the sampling strategy or the way we might represent a sequence of samples (e.g., the model of nested support), the number of unique sub-groups present will grow with sample size. Crucially for scaling laws, the *nature* of that growth—and thus the particular composition of sub-groups and communities contained in the evaluation dataset—is heavily impacted by the sampling strategy used to collect the evaluation dataset.

Sub-groups Can Have Incompatible Metrics

Along with the substantial heterogeneity of populations comes heterogeneity of preferences and values. Although different sub-groups may have different relationships with a single evaluation metric (as in dis-aggregated evaluations of algorithmic fairness (Barocas et al. 2021; Rolf et al. 2022; Burnell et al. 2023)), we are particularly interested in different, incompatible metrics and constructs themselves (Green 2021; Jacobs and Wallach 2021), a subtle but important difference. For instance, it is well-established via decades of the World Values Survey that there are tensions in values across international populations (Haerpfer et al. 2020). In addition, there is large sub-national variation in public opinion; for example, in the US, public opinion differs greatly on topics such as support for gay marriage (Caughey and Warshaw 2019), the New Deal (Caughey, Dougal, and Schickler 2018), and the death penalty (Shirley and Gelman 2015), among others (Berman and Orttung 2020), in ways that are shaped by various cultural and political factors (Caughey and Warshaw 2019).

In AI ethics, prior work has identified substantial differences in how various populations’ values manifest in terms of preferences for AI systems. For instance, Jobin, Ienca, and Vayena (2019) analyzed AI ethics principles statements from nearly a hundred different institutions across various countries, finding that while there was convergence in high-level values such as fairness and transparency, there was high divergence between countries in the specific ways those values are operationalized in AI principles statements, the practices that enact those principles, and the mechanisms used to enforce them. In addition, Awad et al. (2018) collected data on millions of people’s preferences for which of two personas an autonomous vehicle should kill in a car accident, via an online tool used in 233 countries and territories. They found substantial cross-cultural variation in preferences (Awad et al. 2018), and they attempted to explain that variation by drawing on various economic and cultural indicators, such as the World Values Survey (Inglehart et al. 2000), Gini coefficient scores (Dorfman 1979), and other cultural frameworks (McSweeney 2002; Hofstede 2011). Relatedly, Jakesch et al. (2022) conducted a survey of how different groups prioritize ethical values in AI development, finding statistically significant differences in how members of different occupations and demographic groups prioritize values such as fairness, privacy, and trans-

parency in particular AI deployment scenarios. Recent work has explored the relationship between different groups' responses to public opinion polls (e.g., Pew American Trends and the World Values Survey) and the output of large language models, finding that large models' output is more similar to the *average* responses from survey respondents in the USA, Canada, and Australia, compared to respondents from other countries (Durmus et al. 2023) and within the US, language models' output reflects certain groups' opinion poll responses more often (Santurkar et al. 2023).

As one example of how cultural differences in values may impact AI development and evaluation, Sambasivan et al. (2021) has identified how algorithmic (un)fairness in the Indian context operates along different axes than those identified in Western contexts. For instance, they found that algorithmic fairness in India entails different sets of sub-groups, frameworks, and methods, including how algorithmic harms are shaped by the forces of caste and religion, among others. They discuss how popular fairness measurements are informed by specific cultural and historical circumstances, such as approaches for measuring disparate impact or disparate treatment arising from US anti-discrimination law (Watkins, McKenna, and Chen 2022), and by Western philosophical frameworks and approaches to justice more generally (Sambasivan et al. 2021). This suggests that large-scale datasets (which may contain numerous communities or sub-groups) may thus inadvertently collapse meaningful differences in those sub-groups' preferences for how values in AI are operationalized—leading to what some have referred to as “aggregation bias” (Suresh and Guttag 2021). There is substantial empirical evidence for how such aggregate approaches to evaluating models may hide disparities between (or within) subpopulations for a fixed evaluation metric. To uncover these disparities, “dis-aggregated evaluations” (Barocas et al. 2021) of model performance are conducted by dis-aggregating a single performance metric across multiple groups. Such approaches have been the foundation of high-profile evaluations of machine learning failures, such as evaluations of how gender recognition systems are less accurate for women of color than others (Buolamwini and Gebru 2018), or how speech recognition systems have higher word error rates for speakers of African-American Language (Koenecke et al. 2020), among other examples (Obermeyer et al. 2019; De Vries et al. 2019; Ngan et al. 2020).

However, prior empirical work found that when AI product teams deployed models “at scale” (i.e., across numerous geographic and cultural contexts), ambiguity about precisely *how* to dis-aggregate evaluations of model performance—in terms of which evaluation metric to use, or along which demographic dimensions to dis-aggregate—posed major obstacles to AI teams' ability to effectively conduct fairness evaluations at scale (Madaio et al. 2022). Indeed, concerns with aggregation bias are only amplified when we move from fixed evaluation metrics to subpopulation-specific metrics, which may be in tension with each other. Further contributing to potential aggregation harms, during the data annotation process, prior work has found substantial disagreement between annotators from different demographic

groups when determining what constitutes hate speech or toxic language (Sap et al. 2019; Thiago, Marcelo, and Gomes 2021; Davidson et al. 2017). Meanwhile, other work suggests that it is crucial to understand the subjective identities of crowdworkers (Díaz et al. 2022; Denton et al. 2021), developing methods for handling disagreement between groups of annotators in situations where there may not be a single “ground truth” in annotation labels (Davani, Díaz, and Prabhakaran 2022; Gordon et al. 2022, 2021).

Scaling laws, which are aggregate evaluations of models' performance across the entire evaluation dataset U , may similarly hide failures or inverse relationships amongst constructs and values when evaluated with different sub-groups contained within the evaluation data.

Risk of Metric Incompatibility Grows With Data Size

Given that the number of sub-groups within the evaluation dataset grows with the size of that data, and these groups may have incompatible values (i.e., constructs for model output quality) and relationships to performance metrics that operationalize those constructs, we turn to our final claim: that the risk of failures or harms of AI systems grows as data size grows. We ground this claim in the design and evaluation of AI, but similar observations have been made in the political science literature relating the size of a polity and an increase in value friction (Dahl and Tufte 1973).

Goodhart's Law suggests that using a proxy metric can lead to over-optimization and degradation of the actual performance on the construct, something confirmed in the AI alignment literature (Gao, Schulman, and Hilton 2022; Skalse et al. 2022). Similarly, using a single dominant metric can lead to over-optimization and degradation of the actual performance for other constructs and values of different communities, especially because they are more likely to be incompatible with the dominant construct. Indeed, work in multi-task learning has demonstrated that, when optimizing for one task, more data can degrade performance for other tasks (Aribandi et al. 2022). In addition, Solaiman and Dennison (2021) find evidence for a scaling law between model size and toxicity—that is, as model size increases, the models were *more* likely to generate toxic language. Similarly, Lin, Hilton, and Evans (2021) found evidence for an inverse scaling law for model size and truthfulness in a question-answering (QA) task (i.e., models were less truthful the larger they were), and Parrish et al. (2021) found that larger models performed worse on the task of detecting biased language, using a bias benchmark dataset they developed for QA. This phenomena has also been shown as the training dataset size is increased, in addition to the model size. When analyzing the LAION datasets for the presence of hateful content in images and alt-text, Birhane et al. (2024) found that as the dataset size increased, the likelihood for models trained on those datasets to label images of Black people's faces as criminals also increased.

Since the number of distinct sub-groups (and thus their respective latent constructs for model quality) represented in evaluation data is likely to grow with dataset size, there is an increasing chance of a dramatically misaligned evaluation

of model quality—leading to potential impacts or harms for communities whose values are not represented by the dominant performance metric used for model evaluation (Ovalle et al. 2023a; Dev et al. 2021; Felkner et al. 2022).

Given that we might observe a disparity in ‘true performance’ across populations in the evaluation data or, more generally, in the target population, we need to quantify the severity of this disparity. The systematic under-performance and exclusion of values from some sub-groups in scaling law analyses raises issues of (un)fairness and justice. While we emphasize that our claims are different from those in the existing algorithmic bias literature that evaluate a fixed metric for different populations (Barocas et al. 2021; Buolamwini and Gebru 2018)—as we are interested here in values tensions that might result in different communities valuing different constructs entirely or different proxy measures to capture those constructs—we can still adopt methods from that community to measure disparity (Barocas, Hardt, and Narayanan 2019). From the perspective of robustness or Rawlsian fairness, we can look at the worst case true performance of a system on a sub-group in terms of its own values (Hashimoto et al. 2018; Kearns et al. 2018; Ghosh, Genuit, and Reagan 2021). Poor performance is likely to be amplified by the fact that the worst off sub-groups are likely to be from groups that historically have not been considered, represented, or participated in AI development processes (Queerinaï et al. 2023). Given the catastrophic deterioration of performance according to Goodhart’s law, other notions of unfairness (in terms of incompatibility of values and the metrics used to operationalize those values) are also more likely to occur as more sub-groups manifest in evaluation data and model performance according to a target metric grows.

Discussion

Implications for Existing Approaches to Scaling Laws

Proponents of scaling laws for AI systems argue for the existence of a power law relationship between the size of a model (i.e., number of parameters, dataset size, compute) and its performance (along some metric). While this narrative has led to increased investment in collection of large datasets (Luccioni and Viviano 2021; Birhane, Prabhu, and Kahembwe 2021) and in ever-larger models and compute power—along with supporting a narrative of progress akin to Moore’s Law—recent work has demonstrated that scaling laws may not hold for particular tasks (Wei, Tay, and Le 2022; McKenzie et al. 2022; Caballero et al. 2022; Lin, Hilton, and Evans 2021; Parrish et al. 2021).

However, while the current work on exploring the limitations of scaling laws (e.g., inverse scaling laws (McKenzie et al. 2022)) has largely kept the same parameters—the relationship between some aspect of model size and model performance—and just uses a different *task* (e.g., generating truthful text, negation, bias detection), we argue that there may be other relevant dimensions along which scaling laws may not hold. For instance, we argue that scaling laws should consider the size of the evaluation dataset

(in addition to the size of the training dataset). In addition, given our argument that different communities represented within a dataset (or impacted by a particular system) may have fundamentally different values and metrics, what might it look like to evaluate scaling laws where the vertical axis, instead of decontextualized model performance metrics like accuracy (or F1, RMSE, ROUGE), were instead chosen for particular use cases or system deployment contexts, or were chosen by particular impacted communities in participatory ways (Wagner et al. 2021; Delgado et al. 2023; Suresh et al. 2022; Dennler et al. 2023) to better reflect their values.

Returning to our introductory example, although mitigations for the curse of multilinguality have been developed (Pfeiffer et al. 2022; Blevins et al. 2024), they require a well-defined set of languages and metrics for each, both of which are missing for the multiplicity of sub-groups we suggest.

Moreover, substantial work on model evaluations has shown that aggregate metrics of model performance may hide worse performance for particular sub-groups that can be observed when model performance is *dis*-aggregated by some demographic categories (Buolamwini and Gebru 2018; Koenecke et al. 2020; Obermeyer et al. 2019; Ngan et al. 2020; Barocas et al. 2021). Analogously, the current paradigm of evaluating scaling laws on aggregations of performance metrics evaluated on a single training dataset is likely to hide similar divergences in values and preferences for metrics for sub-groups within an evaluation dataset. For instance, a recent paper proposed a benchmark for evaluating bias in QA, and evaluated it on several large language models; however, they caveat at the end of the paper that “*the data in BBQ is only designed to test biases associated with US English-speaking cultural contexts, and it should not be used as evidence that a model would still look unbiased for contexts from a different culture*” (Parrish et al. 2021). What would it look like for evaluations of scaling laws to be dis-aggregated for datasets collected from (or ideally collected, curated, or annotated by) different communities (Wagner et al. 2021), be that linguistic communities, cultural communities, countries, or other sets of sub-groups? In other words, would an observed scaling law relationship for performance metric μ evaluated with dataset U still hold if that dataset were collected from a different context or collected or annotated by a different community?

Broader Questions for Scaling Laws

While the previous section suggests some shorter-term means to investigate the limits of scaling laws, this work raises more fundamental questions for scaling laws for AI. For instance, while analyses like the ones we proposed may reveal broken, inverse, or other non-monotonically increasing functional forms for scaling laws across different communities within a given dataset (or who might use or be impacted by a particular system), what to do about that is a much thornier question. In many ways, tensions in values (sometimes referred to as values pluralism (Berlin 1969; Crowder 2002; Van der Wal and Van Hout 2009)) are a fundamental challenge of political systems—including technology, as technologies enact politics (Winner 1987; Mulligan and Bamberger 2019; Nissenbaum 2011). Meth-

ods have been developed from participatory democracy (Polletta 2012), deliberation theory (Fishkin and Luskin 2005), and value-sensitive design (Miller et al. 2007) (among other areas) to identify and resolve tensions in values. However, these methods have largely been designed for smaller scales: e.g., town halls, focus group discussions, and participatory design workshops (Muller and Druin 2012). As such, it is not clear how such approaches may be able to address value tensions at the scale of modern AI systems (Sloane et al. 2020; Hanna and Park 2020; Young et al. 2024). This may be an argument for drawing on related work from anthropology that argues for “non-scalability” as a desired quality of social systems that resist the impetus to scale (Sharma, Kumar, and Nardi 2023), or that rethink the nature of systems’ design as they move from context to context (Tsing 2012), rather than relying on the supposed portability of sociotechnical systems (Selbst et al. 2019).

Recent work on AI “alignment” has attempted to develop approaches to aligning AI with human values (Askell et al. 2021; Bai et al. 2022; Ganguli et al. 2022b; Solaiman and Dennison 2021). Gabriel (2020) discusses how values pluralism may impact the goal of aligning AI systems with “human values” (in all of their multiplicity and tensions), and he discusses tradeoffs in several potential approaches to resolving those tensions. Some currently adopted approaches include red-teaming (Ganguli et al. 2022b), reinforcement learning from human feedback (RLHF) (Bai et al. 2022) or creating “values-targeted datasets” (Solaiman and Dennison 2021). However, in recent work attempting to incorporate values into training large models, people involved in red-teaming and RLHF appear to largely be US-based and may not be representative even of the US population (Ganguli et al. 2022b; Bai et al. 2022). Meanwhile, in this example, the “values-targeted datasets” were created by the researchers, who acknowledge this limitation, writing that “*creating many values-targeted datasets to reflect the cultures of the many peoples impacted by language models is a difficult feat*” (Solaiman and Dennison 2021).

This is not, however, simply a limitation of their work, but a more fundamental challenge for values tensions in scaling laws—what might it mean to not only create values-targeted evaluation datasets from different communities or cultural contexts, but to resolve potentially irreconcilable differences in values between such communities, at scale? Is such a goal itself misguided? While we proposed a simple worst-case approach to quantify the robustness of a model across subpopulations, the reality of dealing with multiple, conflicting values is more complex. The evidence from Birhane et al. (2024) demonstrates that values of emergent subpopulations can be toxic, suggesting that simply looking at subpopulations without context risks buttressing toxic behavior.

Although we do not offer answers to the questions we have proposed in this paper, we suggest that, in part, what is needed are interdisciplinary, mixed-methods approaches to theoretically and empirically investigate the questions we raise. There are existing theories and methods from numerous fields—across the social sciences and computer science—that have been developed to explore questions related to whether and in what ways various communities or

sub-groups are represented by data (Yoo 2015) as well as identifying and resolving values tensions among communities (Miller et al. 2007). For instance, various methods have been developed for identifying sub-groups within large datasets about people, in political science (e.g., for public opinion polling) (Feczko et al. 2019), demography (Page 2014), healthcare (Wick et al. 2022), genetics (Patterson, Price, and Reich 2006), and more. However, work from the social sciences suggests that some communities may be hidden in ways not legible to data available for computational community detection (e.g., injection drug users; sex workers) (Salganik and Heckathorn 2004), requiring approaches such as qualitative or ethnographic research to identify such communities—which may be difficult and prohibitively expensive in current practice at the scales implicated by AI scaling laws, and moreover, many vulnerable or marginalized communities may not want to have data collected at all, as it may put them at risk (Muller and Strohmayer 2022; Queerina et al. 2023; Dennler et al. 2023).

In addition, recent work demonstrating the lack of relevance of Western AI fairness frameworks for other cultural contexts, such as India, has drawn on approaches from both qualitative (Sambasivan et al. 2021) and quantitative research paradigms, including natural language processing (Bhatt et al. 2022). We argue that interdisciplinary, mixed-methods approaches such as these, involving deep partnerships with community members or community organizations, such as participatory or community-based research methods (Delgado et al. 2023; Wagner et al. 2021; Sambasivan et al. 2021; Harrington, Erete, and Piper 2019; Muller and Druin 2012; Dennler et al. 2023) may be one way to empirically investigate our claims and grapple with the inherent tensions and limitations of scaling laws for AI.

Conclusion

Analyses of AI scaling laws paint an impressive picture of progress in AI driven by increasing scales of data, model size, and compute power. We argue that these results are not the *whole* picture. We draw together work from computer science and the social sciences to identify the ways that metrics used to evaluate AI systems may be unstable or precarious; how the increasing scales of data used to train AI systems entail increasing numbers of sub-groups or communities of people; and how those groups’ values and preferences for AI systems and the metrics used to evaluate them may be incompatible or fundamentally at odds with each other. We close by discussing how these insights pose fundamental challenges to the paradigm of AI scaling laws, and we raise open questions and opportunities for the field to investigate whether, when, and for whom scaling laws may (or may not) hold. We suggest opportunities for interdisciplinary, participatory, and community-based research to better understand which sub-groups or communities may be represented in a given dataset (or impacted by a particular model); which evaluation metrics might best reflect their values; and how to conduct such evaluations or resolve tensions in those values. The results of this work may suggest that progress in large AI models is not quite as straightforward as it first appears. In other words, one might ask: progress for whom?

References

- Aghajanyan, A.; Yu, L.; Conneau, A.; Hsu, W.-N.; Hambarzumyan, K.; Zhang, S.; Roller, S.; Goyal, N.; Levy, O.; and Zettlemoyer, L. 2023. Scaling laws for generative mixed-modal language models. In *Proceedings of the 40th International Conference on Machine Learning, ICML'23*. JMLR.org.
- Aker, J. C.; and Mbiti, I. M. 2010. Mobile phones and economic development in Africa. *Journal of economic Perspectives*, 24(3): 207–232.
- Altman, S. 2023. Planning for AGI and beyond. *OpenAI Blog*, February.
- Anderson, C. 2008. The end of theory: The data deluge makes the scientific method obsolete. *Wired magazine*, 16(7): 16–07.
- Aribandi, V.; Tay, Y.; Schuster, T.; Rao, J.; Zheng, H. S.; Mehta, S. V.; Zhuang, H.; Tran, V. Q.; Bahri, D.; Ni, J.; Gupta, J.; Hui, K.; Ruder, S.; and Metzler, D. 2022. ExT5: Towards Extreme Multi-Task Scaling for Transfer Learning. In *International Conference on Learning Representations*.
- Askell, A.; Bai, Y.; Chen, A.; Drain, D.; Ganguli, D.; Henighan, T.; Jones, A.; Joseph, N.; Mann, B.; DasSarma, N.; et al. 2021. A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861*.
- Awad, E.; Dsouza, S.; Kim, R.; Schulz, J.; Henrich, J.; Shariff, A.; Bonnefon, J.-F.; and Rahwan, I. 2018. The moral machine experiment. *Nature*, 563(7729): 59–64.
- Bai, Y.; Jones, A.; Ndousse, K.; Askell, A.; Chen, A.; DasSarma, N.; Drain, D.; Fort, S.; Ganguli, D.; Henighan, T.; et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Barocas, S.; Guo, A.; Kamar, E.; Krones, J.; Morris, M. R.; Vaughan, J. W.; Wadsworth, W. D.; and Wallach, H. 2021. Designing disaggregated evaluations of ai systems: Choices, considerations, and tradeoffs. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 368–378.
- Barocas, S.; Hardt, M.; and Narayanan, A. 2019. *Fairness and Machine Learning: Limitations and Opportunities*. fairmlbook.org. <http://www.fairmlbook.org>.
- Beer, D. 2016. *Metric power*. Springer.
- Bender, E. M.; Gebru, T.; McMillan-Major, A.; and Shmitchell, S. 2021. On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 610–623.
- Bergman, A. S.; Hendricks, L. A.; Rauh, M.; Wu, B.; Agnew, W.; Kunesch, M.; Duan, I.; Gabriel, I.; and Isaac, W. 2023. Representation in AI Evaluations. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 519–533.
- Berlin, I. 1969. Four essays on liberty.
- Berman, M.; and Orttung, R. W. 2020. Measuring progress toward urban sustainability: Do global measures work for Arctic cities? *Sustainability*, 12(9): 3708.
- Bernstein, E. S. 2017. Making transparency transparent: The evolution of observation in management theory. *Academy of Management Annals*, 11(1): 217–266.
- Bhatt, S.; Dev, S.; Talukdar, P.; Dave, S.; and Prabhakaran, V. 2022. Cultural Re-contextualization of Fairness Research in Language Technologies in India. *arXiv preprint arXiv:2211.11206*.
- Birhane, A.; Dehdashtian, S.; Prabhu, V. U.; and Boddeti, V. 2024. The Dark Side of Dataset Scaling: Evaluating Racial Classification in Multimodal Models. In *2024 ACM Conference on Fairness, Accountability, and Transparency, FAccT '24*. Association for Computing Machinery.
- Birhane, A.; Kalluri, P.; Card, D.; Agnew, W.; Dotan, R.; and Bao, M. 2022. The Values Encoded in Machine Learning Research. In *2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22*, 173–184. New York, NY, USA: Association for Computing Machinery.
- Birhane, A.; Prabhu, V. U.; and Kahembwe, E. 2021. Multimodal datasets: misogyny, pornography, and malignant stereotypes. *arXiv preprint arXiv:2110.01963*.
- Blevins, T.; Limisiewicz, T.; Gururangan, S.; Li, M.; Gonen, H.; Smith, N. A.; and Zettlemoyer, L. 2024. Breaking the Curse of Multilinguality with Cross-lingual Expert Language Models. *arXiv:2401.10440*.
- boyd, d.; and Crawford, K. 2012. Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, communication & society*, 15(5): 662–679.
- Broder, A. 2002. A taxonomy of web search. *SIGIR Forum*, 36(2): 3–10.
- Buolamwini, J.; and Gebru, T. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, 77–91. PMLR.
- Burnell, R.; Schellaert, W.; Burden, J.; Ullman, T. D.; Martinez-Plumed, F.; Tenenbaum, J. B.; Rutar, D.; Cheke, L. G.; Sohl-Dickstein, J.; Mitchell, M.; Kiela, D.; Shanahan, M.; Voorhees, E. M.; Cohn, A. G.; Leibo, J. Z.; and Hernandez-Orallo, J. 2023. Rethink reporting of evaluation results in AI. *Science*, 380(6641): 136–138.
- Caballero, E.; Gupta, K.; Rish, I.; and Krueger, D. 2022. Broken Neural Scaling Laws.
- Campbell, D. T. 1979. Assessing the impact of planned social change. *Evaluation and program planning*, 2(1): 67–90.
- Carterette, B.; Kanoulas, E.; and Yilmaz, E. 2012. Advances on the Development of Evaluation Measures. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '12*, 1200–1201. New York, NY, USA: ACM.
- Caughey, D.; Dougal, M.; and Schickler, E. 2018. The Policy Bases of the New Deal Realignment: Evidence from Public Opinion Polls, 1936–1952. *Journal of Politics*.
- Caughey, D.; and Warshaw, C. 2019. Public Opinion in Subnational Politics. *Journal of Politics*, 81(1).
- Chandar, P.; Diaz, F.; and St. Thomas, B. 2020a. Beyond Accuracy: Grounding Evaluation Metrics for Human-Machine

- Learning Systems. In *Advances in Neural Information Processing Systems*.
- Chandar, P.; Diaz, F.; and St. Thomas, B. 2020b. Beyond Accuracy: Grounding Evaluation Metrics for Human-Machine Learning Systems. In *Advances in Neural Information Processing Systems*.
- Chang, T. A.; Arnett, C.; Tu, Z.; and Bergen, B. K. 2023. When is Multilinguality a Curse? Language Modeling for 250 High- and Low-Resource Languages. *arXiv preprint*.
- Chasalow, K.; and Levy, K. 2021. Representativeness in statistics, politics, and machine learning. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 77–89.
- Chen, Y.; Conroy, N. J.; and Rubin, V. L. 2015. Misleading Online Content: Recognizing Clickbait as “False News”. In *Proceedings of the 2015 ACM on Workshop on Multimodal Deception Detection, WMDD ’15*, 15–19. New York, NY, USA: Association for Computing Machinery. ISBN 9781450339872.
- Cherti, M.; Beaumont, R.; Wightman, R.; Wortsman, M.; Ilharco, G.; Gordon, C.; Schuhmann, C.; Schmidt, L.; and Jitsev, J. 2022. Reproducible scaling laws for contrastive language-image learning.
- Chouldechova, A. 2017. Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. *Big Data*, 5(2): 153–163. PMID: 28632438.
- Chowdhery, A.; Narang, S.; Devlin, J.; Boma, M.; Mishra, G.; Roberts, A.; Barham, P.; Chung, H. W.; Sutton, C.; Gehrmann, S.; Schuh, P.; Shi, K.; Tsvyashchenko, S.; Maynez, J.; Rao, A.; Barnes, P.; Tay, Y.; Shazeer, N.; Prabhakaran, V.; Reif, E.; Du, N.; Hutchinson, B.; Pope, R.; Bradbury, J.; Austin, J.; Isard, M.; Gur-Ari, G.; Yin, P.; Duke, T.; Levskaya, A.; Ghemawat, S.; Dev, S.; Michalewski, H.; Garcia, X.; Misra, V.; Robinson, K.; Fedus, L.; Zhou, D.; Ippolito, D.; Luan, D.; Lim, H.; Zoph, B.; Spiridonov, A.; Sepassi, R.; Dohan, D.; Agrawal, S.; Omer-nick, M.; Dai, A. M.; Pillai, T. S.; Pellat, M.; Lewkowycz, A.; Moreira, E.; Child, R.; Polozov, O.; Lee, K.; Zhou, Z.; Wang, X.; Saeta, B.; Diaz, M.; Firat, O.; Catasta, M.; Wei, J.; Meier-Hellstern, K.; Eck, D.; Dean, J.; Petrov, S.; and Fiedel, N. 2022. PaLM: Scaling Language Modeling with Pathways.
- Christin, A. 2020. *Metrics at Work: Journalism and the Contested Meaning of Algorithms*. Princeton University Press.
- Cohen, P. R. 1995. *Empirical methods for artificial intelligence*. MIT Press.
- Cohen, P. R.; and Howe, A. E. 1988. How Evaluation Guides AI Research. *AI Magazine*, 9(4): 35–43.
- Collins, P. H. 2019. *Intersectionality as critical social theory*. Duke University Press.
- Conneau, A.; Khandelwal, K.; Goyal, N.; Chaudhary, V.; Wenzek, G.; Guzmán, F.; Grave, E.; Ott, M.; Zettlemoyer, L.; and Stoyanov, V. 2020. Unsupervised Cross-lingual Representation Learning at Scale. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 8440–8451. Online: Association for Computational Linguistics.
- Crenshaw, K. 1990. Mapping the margins: Intersectionality, identity politics, and violence against women of color. *Stan. L. Rev.*, 43: 1241.
- Crowder, G. 2002. *Liberalism and value pluralism*. Bloomsbury Publishing.
- Dahl, R. A.; and Tufte, E. R. 1973. *Size and Democracy*. Stanford University Press.
- Davani, A. M.; Díaz, M.; and Prabhakaran, V. 2022. Dealing with disagreements: Looking beyond the majority vote in subjective annotations. *Transactions of the Association for Computational Linguistics*, 10: 92–110.
- Davidson, T.; Warmley, D.; Macy, M.; and Weber, I. 2017. Automated hate speech detection and the problem of offensive language. In *Proceedings of the international AAAI conference on web and social media*, volume 11, 512–515.
- de Marchi, S.; and Page, S. E. 2016. Computational Social Science: Discovery and Prediction. Edited by R. Michael Alvarez. New York: Cambridge University Press, 2016. 337p. 34.99 paper. *Perspectives on Politics*, 14(4): 1169–1170.
- De Vries, T.; Misra, I.; Wang, C.; and Van der Maaten, L. 2019. Does object recognition work for everyone? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 52–59.
- Delgado, F.; Yang, S.; Madaio, M.; and Yang, Q. 2023. The participatory turn in ai design: Theoretical foundations and the current state of practice. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–23.
- Dennler, N.; Ovalle, A.; Singh, A.; Soldaini, L.; Subramonian, A.; Tu, H.; Agnew, W.; Ghosh, A.; Yee, K.; Peradejordi, I. F.; et al. 2023. Bound by the Bounty: Collaboratively Shaping Evaluation Processes for Queer AI Harms. *arXiv preprint arXiv:2307.10223*.
- Denton, E.; Díaz, M.; Kivlichan, I.; Prabhakaran, V.; and Rosen, R. 2021. Whose ground truth? accounting for individual and collective identities underlying dataset annotation. *arXiv preprint arXiv:2112.04554*.
- Dettmers, T.; and Zettlemoyer, L. 2022. The case for 4-bit precision: k-bit Inference Scaling Laws.
- Dev, S.; Monajatipoor, M.; Ovalle, A.; Subramonian, A.; Phillips, J. M.; and Chang, K.-W. 2021. Harms of gender exclusivity and challenges in non-binary representation in language technologies. *arXiv preprint arXiv:2108.12084*.
- Díaz, M.; Kivlichan, I.; Rosen, R.; Baker, D.; Amironesei, R.; Prabhakaran, V.; and Denton, E. 2022. Crowdwork-sheets: Accounting for individual and collective identities underlying crowdsourced dataset annotation. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, 2342–2351.
- DiSalvo, C. 2009. Design and the Construction of Publics. *Design issues*, 25(1): 48–63.

- Dodge, J.; Sap, M.; Marasović, A.; Agnew, W.; Ilharco, G.; Groeneveld, D.; Mitchell, M.; and Gardner, M. 2021. Documenting large webtext corpora: A case study on the colossal clean crawled corpus. *arXiv preprint arXiv:2104.08758*.
- Dorfman, R. 1979. A formula for the Gini coefficient. *The review of economics and statistics*, 146–149.
- Durmus, E.; Nyugen, K.; Liao, T. I.; Schiefer, N.; Askill, A.; Bakhtin, A.; Chen, C.; Hatfield-Dodds, Z.; Hernandez, D.; Joseph, N.; et al. 2023. Towards Measuring the Representation of Subjective Global Opinions in Language Models. *arXiv preprint arXiv:2306.16388*.
- Feczko, E.; Miranda-Dominguez, O.; Marr, M.; Graham, A. M.; Nigg, J. T.; and Fair, D. A. 2019. The heterogeneity problem: approaches to identify psychiatric subtypes. *Trends in cognitive sciences*, 23(7): 584–601.
- Felkner, V. K.; Chang, H.-C. H.; Jang, E.; and May, J. 2022. Towards WinoQueer: Developing a Benchmark for Anti-Queer Bias in Large Language Models. *arXiv preprint arXiv:2206.11484*.
- Fishkin, J. S.; and Luskin, R. C. 2005. Experimenting with a democratic ideal: Deliberative polling and public opinion. *Acta politica*, 40: 284–298.
- Friedler, S. A.; Scheidegger, C.; and Venkatasubramanian, S. 2016. On the (im)possibility of fairness. *CoRR*, abs/1609.07236.
- Gabriel, I. 2020. Artificial Intelligence, Values, and Alignment. *Minds and Machines*, 30(3): 411–437.
- Gabriel, T. 2010. Under pressure, teachers tamper with tests. *The New York Times*, 10.
- Ganesan, K. 2018. Rouge 2.0: Updated and improved measures for evaluation of summarization tasks. *arXiv preprint arXiv:1803.01937*.
- Ganguli, D.; Hernandez, D.; Lovitt, L.; Askill, A.; Bai, Y.; Chen, A.; Conerly, T.; Dassarma, N.; Drain, D.; Elhage, N.; El Showk, S.; Fort, S.; Hatfield-Dodds, Z.; Henighan, T.; Johnston, S.; Jones, A.; Joseph, N.; Kernian, J.; Kravec, S.; Mann, B.; Nanda, N.; Ndousse, K.; Olsson, C.; Amodei, D.; Brown, T.; Kaplan, J.; McCandlish, S.; Olah, C.; Amodei, D.; and Clark, J. 2022a. Predictability and Surprise in Large Generative Models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, 1747–1764. New York, NY, USA: Association for Computing Machinery.
- Ganguli, D.; Lovitt, L.; Kernion, J.; Askill, A.; Bai, Y.; Kadavath, S.; Mann, B.; Perez, E.; Schiefer, N.; Ndousse, K.; et al. 2022b. Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned. *arXiv preprint arXiv:2209.07858*.
- Gao, L.; Schulman, J.; and Hilton, J. 2022. Scaling Laws for Reward Model Overoptimization.
- Ghassemi, M.; and Mohamed, S. 2022. Machine learning and health need better values. *NPJ Digital Medicine*, 5(1): 51.
- Ghosh, A.; Genuit, L.; and Reagan, M. 2021. Characterizing Intersectional Group Fairness with Worst-Case Comparisons. In Lamba, D.; and Hsu, W. H., eds., *Proceedings of 2nd Workshop on Diversity in Artificial Intelligence (AID-BEI)*, volume 142 of *Proceedings of Machine Learning Research*, 22–34. PMLR.
- Goodhart, C. A. 1984. *Problems of monetary management: the UK experience*. Springer.
- Gordon, M. L.; Lam, M. S.; Park, J. S.; Patel, K.; Hancock, J.; Hashimoto, T.; and Bernstein, M. S. 2022. Jury Learning: Integrating Dissenting Voices into Machine Learning Models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22. New York, NY, USA: Association for Computing Machinery.
- Gordon, M. L.; Zhou, K.; Patel, K.; Hashimoto, T.; and Bernstein, M. S. 2021. The Disagreement Deconvolution: Bringing Machine Learning Performance Metrics In Line With Reality. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21. New York, NY, USA: Association for Computing Machinery.
- Gray, D.; Micheli, P.; and Pavlov, A. 2015. *Measurement Madness: Recognizing and avoiding the pitfalls of performance measurement*. John Wiley & Sons.
- Gray, M. L.; and Suri, S. 2019. *Ghost work: How to stop Silicon Valley from building a new global underclass*. Eamon Dolan Books.
- Green, B. 2019. Good” isn’t good enough. In *Proceedings of the AI for Social Good workshop at NeurIPS*, volume 17.
- Green, B. 2021. The contestation of tech ethics: A sociotechnical approach to technology ethics in practice. *Journal of Social Computing*, 2(3): 209–225.
- Griesemer, J. 2020. Taking Goodhart’s law meta: Gaming, meta-gaming, and hacking academic performance metrics. *Gaming the Metrics*, 77.
- Grotov, A.; and de Rijke, M. 2016. Online Learning to Rank for Information Retrieval: SIGIR 2016 Tutorial. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '16, 1215–1218. New York, NY, USA: Association for Computing Machinery.
- Haerpfner, C.; Inglehart, R.; Moreno, A.; Welzel, C.; Kizilova, K.; Diez-Medrano, J.; Lagos, M.; Norris, P.; Ponarin, E.; Puranen, B.; et al. 2020. World values survey: round seven–country-pooled datafile. *Madrid, Spain & Vienna, Austria: JD Systems Institute & WVSA Secretariat*.
- Hand, D. J. 2010. *Measurement Theory and Practice: The World Through Quantification*. Wiley.
- Hanna, A.; and Park, T. M. 2020. Against Scale: Provocations and Resistances to Scale Thinking.
- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29.
- Hargittai, E. 2015. Is bigger always better? Potential biases of big data derived from social network sites. *The ANNALS of the American Academy of Political and Social Science*, 659(1): 63–76.
- Hargittai, E. 2020. Potential biases in big data: Omitted voices on social media. *Social Science Computer Review*, 38(1): 10–24.

- Harrington, C.; Erete, S.; and Piper, A. M. 2019. Deconstructing community-based collaborative design: Towards more equitable participatory design engagements. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW): 1–25.
- Hashimoto, T.; Srivastava, M.; Namkoong, H.; and Liang, P. 2018. Fairness Without Demographics in Repeated Loss Minimization. In Dy, J.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 1929–1938. Stockholmsmässan, Stockholm Sweden: PMLR.
- Hilton, J.; Tang, J.; and Schulman, J. 2023. Scaling laws for single-agent reinforcement learning.
- Hofstede, G. 2011. Dimensionalizing Cultures: The Hofstede Model in Context. *Online readings in psychology and culture*, 2(1): 2307–0919.
- Inglehart, R.; Basanez, M.; Diez-Medrano, J.; Halman, L.; and Luijckx, R. 2000. World values surveys and European values surveys, 1981–1984, 1990–1993, and 1995–1997. *Ann Arbor-Michigan, Institute for Social Research, ICPSR version*.
- Ivgi, M.; Carmon, Y.; and Berant, J. 2022. Scaling Laws Under the Microscope: Predicting Transformer Performance from Small Scale Experiments.
- Jacobs, A. Z.; and Wallach, H. 2021. Measurement and Fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’21, 375–385. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383097.
- Jakesch, M.; Buçinca, Z.; Amershi, S.; and Olteanu, A. 2022. How different groups prioritize ethical values for responsible AI. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, 310–323.
- Joachims, T. 2002. Optimizing search engines using click-through data. In *KDD ’02: Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, 133–142.
- Jobin, A.; Ienca, M.; and Vayena, E. 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9): 389–399.
- Kamvar, M.; and Baluja, S. 2006. A Large Scale Study of Wireless Search Behavior: Google Mobile Search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’06, 701–709. New York, NY, USA: Association for Computing Machinery. ISBN 1595933727.
- Kamvar, M.; and Baluja, S. 2007. Deciphering Trends in Mobile Search. *Computer*, 40(8): 58–62.
- Kaplan, J.; McCandlish, S.; Henighan, T.; Brown, T. B.; Chess, B.; Child, R.; Gray, S.; Radford, A.; Wu, J.; and Amodei, D. 2020. Scaling Laws for Neural Language Models.
- Kearns, M.; Neel, S.; Roth, A.; and Wu, Z. S. 2018. Preventing Fairness Gerrymandering: Auditing and Learning for Subgroup Fairness. In Dy, J.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 2564–2572. PMLR.
- Khan, M. A. A.; Khan, M. M. R.; Hassan, M.; Ahmed, F.; and Haque, S. M. R. 2017. Role of community radio for community development in Bangladesh. *The International Technology Management Review*, 6(3): 94–102.
- Kitchin, R. 2014. Big Data, new epistemologies and paradigm shifts. *Big data & society*, 1(1): 2053951714528481.
- Koenecke, A.; Nam, A.; Lake, E.; Nudell, J.; Quartey, M.; Mengesha, Z.; Toups, C.; Rickford, J. R.; Jurafsky, D.; and Goel, S. 2020. Racial disparities in automated speech recognition. *Proceedings of the National Academy of Sciences*, 117(14): 7684–7689.
- Kohavi, R.; Tang, D.; and Xu, Y. 2020. *Trustworthy Online Controlled Experiments: A Practical Guide to A/B Testing*. Cambridge University Press.
- Lin, S.; Hilton, J.; and Evans, O. 2021. Truthfulqa: Measuring how models mimic human falsehoods. *arXiv preprint arXiv:2109.07958*.
- Luccioni, A.; and Viviano, J. 2021. What’s in the box? an analysis of undesirable content in the Common Crawl corpus. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, 182–189.
- Madaio, M.; Egede, L.; Subramonyam, H.; Wortman Vaughan, J.; and Wallach, H. 2022. Assessing the Fairness of AI Systems: AI Practitioners’ Processes, Challenges, and Needs for Support. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW1): 1–26.
- Magueresse, A.; Carles, V.; and Heetderks, E. 2020. Low-resource languages: A review of past work and future challenges. *arXiv preprint arXiv:2006.07264*.
- McKenzie, I.; Lyzhov, A.; Parrish, A.; Prabhu, A.; Mueller, A.; Kim, N.; Bowman, S.; and Perez, E. 2022. The Inverse Scaling Prize.
- McSweeney, B. 2002. Hofstede’s model of national cultural differences and their consequences: A triumph of faith—a failure of analysis. *Human relations*, 55(1): 89–118.
- Miller, J. K.; Friedman, B.; Jancke, G.; and Gill, B. 2007. Value tensions in design: the value sensitive design, development, and appropriation of a corporation’s groupware system. In *Proceedings of the 2007 international ACM conference on Supporting group work*, 281–290.
- Mitchell, A.; and Hitlin, P. 2013. Twitter reaction to events often at odds with overall public opinion.
- Mügge, D. 2022. Economic statistics as political artefacts. *Review of international political economy*, 29(1): 1–22.
- Muller, M.; and Strohmayer, A. 2022. Forgetting practices in the data sciences. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–19.

- Muller, M. J.; and Druin, A. 2012. Participatory design: The third space in human-computer interaction. In *The Human-Computer Interaction Handbook*, 1125–1153. CRC Press.
- Mulligan, D. K.; and Bamberger, K. A. 2019. Procurement as policy: Administrative process for machine learning. *Berkeley Tech. LJ*, 34: 773.
- Murdock, V.; and Croft, W. B. 2002. Task orientation in question answering. In *SIGIR '02: Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, 355–356. New York, NY, USA: ACM.
- Ngan, M.; Ngan, M.; Grother, P.; Hanaoka, K.; and Kuo, J. 2020. *Face recognition vendor test (frvt) part 4: Morph-performance of automated face morph detection*. US Department of Commerce, National Institute of Standards and Technology.
- Nissenbaum, H. 2011. From preemption to circumvention: if technology regulates, why do we need regulation (and vice versa). *Berkeley Tech. LJ*, 26: 1367.
- Obermeyer, Z.; Powers, B.; Vogeli, C.; and Mullainathan, S. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464): 447–453.
- Okeleke, K.; and Suardi, S. 2019. The Mobile Economy Sub-Saharan Africa 2021. *GSMA Intelligence Report*.
- Ovalle, A.; Goyal, P.; Dhamala, J.; Jaggers, Z.; Chang, K.-W.; Galstyan, A.; Zemel, R.; and Gupta, R. 2023a. “I’m fully who I am”: Towards Centering Transgender and Non-Binary Voices to Measure Biases in Open Language Generation. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 1246–1266.
- Ovalle, A.; Subramonian, A.; Gautam, V.; Gee, G.; and Chang, K.-W. 2023b. Factoring the Matrix of Domination: A Critical Review and Reimagination of Intersectionality in AI Fairness. *arXiv preprint arXiv:2303.17555*.
- Page, S. E. 2014. Where diversity comes from and why it matters? *European Journal of Social Psychology*, 44(4): 267–279.
- Parrish, A.; Chen, A.; Nangia, N.; Padmakumar, V.; Phang, J.; Thompson, J.; Htut, P. M.; and Bowman, S. R. 2021. BBQ: A hand-built bias benchmark for question answering. *arXiv preprint arXiv:2110.08193*.
- Patterson, N.; Price, A. L.; and Reich, D. 2006. Population structure and eigenanalysis. *PLoS genetics*, 2(12): e190.
- Perrin, A. 2015. Social Media Usage: 2005-2015. Technical report, Pew Research Center.
- Pfeiffer, J.; Goyal, N.; Lin, X.; Li, X.; Cross, J.; Riedel, S.; and Artetxe, M. 2022. Lifting the Curse of Multilinguality by Pre-training Modular Transformers. In Carpuat, M.; de Marneffe, M.-C.; and Meza Ruiz, I. V., eds., *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3479–3495. Seattle, United States: Association for Computational Linguistics.
- Pinch, T. J.; and Bijker, W. E. 1984. The social construction of facts and artefacts: Or how the sociology of science and the sociology of technology might benefit each other. *Social studies of science*, 14(3): 399–441.
- Polletta, F. 2012. Freedom is an endless meeting. In *Freedom Is an Endless Meeting*. University of Chicago Press.
- Post, M. 2018. A call for clarity in reporting BLEU scores. *arXiv preprint arXiv:1804.08771*.
- Poushter, J.; Bishop, C.; and Chwe, H. 2018. Social Media Use Continues to Rise in Developing Countries but Plateaus Across Developed Ones. Technical report, Pew Research Center.
- Queerina, O. O.; Ovalle, A.; Subramonian, A.; Singh, A.; Voelcker, C.; Sutherland, D. J.; Locatelli, D.; Breznik, E.; Klubicka, F.; Yuan, H.; et al. 2023. Queer In AI: A Case Study in Community-Led Participatory AI. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 1882–1895.
- Rae, J. W.; Borgeaud, S.; Cai, T.; Millican, K.; Hoffmann, J.; Song, F.; Aslanides, J.; Henderson, S.; Ring, R.; Young, S.; Rutherford, E.; Hennigan, T.; Menick, J.; Cassirer, A.; Powell, R.; Driessche, G. v. d.; Hendricks, L. A.; Rauh, M.; Huang, P.-S.; Glaese, A.; Welbl, J.; Dathathri, S.; Huang, S.; Uesato, J.; Mellor, J.; Higgins, I.; Creswell, A.; McAleese, N.; Wu, A.; Elsen, E.; Jayakumar, S.; Buchatskaya, E.; Budden, D.; Sutherland, E.; Simonyan, K.; Paganini, M.; Sifre, L.; Martens, L.; Li, X. L.; Kuncoro, A.; Nematzadeh, A.; Gribovskaya, E.; Donato, D.; Lazaridou, A.; Mensch, A.; Lepiau, J.-B.; Tsimpoukelli, M.; Grigorev, N.; Fritz, D.; Sottiaux, T.; Pajarskas, M.; Pohlen, T.; Gong, Z.; Toyama, D.; d’Autume, C. d. M.; Li, Y.; Terzi, T.; Mikulik, V.; Babuschkin, I.; Clark, A.; Casas, D. d. L.; Guy, A.; Jones, C.; Bradbury, J.; Johnson, M.; Hechtman, B.; Weidinger, L.; Gabriel, I.; Isaac, W.; Lockhart, E.; Osindero, S.; Rimell, L.; Dyer, C.; Vinyals, O.; Ayoub, K.; Stanway, J.; Bennett, L.; Hassabis, D.; Kavukcuoglu, K.; and Irving, G. 2021. Scaling Language Models: Methods, Analysis & Insights from Training Gopher.
- Raji, I. D.; Bender, E. M.; Paullada, A.; Denton, E.; and Hanna, A. 2021. AI and the everything in the whole wide world benchmark. *arXiv preprint arXiv:2111.15366*.
- Ranganathan, A.; and Benson, A. 2020. A numbers game: Quantification of work, auto-gamification, and worker productivity. *American Sociological Review*, 85(4): 573–609.
- Rei, R.; Stewart, C.; Farinha, A. C.; and Lavie, A. 2020. COMET: A Neural Framework for MT Evaluation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2685–2702. Online: Association for Computational Linguistics.
- Reiter, E. 2018. A structured review of the validity of BLEU. *Computational Linguistics*, 44(3): 393–401.
- Rish, I. 2023. On Scaling Laws, Emergent Behaviors, and AI Democratization Efforts. Keynote, Predictable AI.
- Robertson, R. E.; Olteanu, A.; Diaz, F.; Shokouhi, M.; and Bailey, P. 2021. “I Can’t Reply with That”: Characterizing Problematic Email Reply Suggestions. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–18.
- Rolf, E.; Packer, B.; Beutel, A.; and Diaz, F. 2022. Striving for data-model efficiency: Identifying data externalities on

- group performance. In *Workshop on Trustworthy and Socially Responsible Machine Learning, NeurIPS 2022*.
- Rosenfeld, J. S.; Rosenfeld, A.; Belinkov, Y.; and Shavit, N. 2020. A Constructive Prediction of the Generalization Error Across Scales. In *International Conference on Learning Representations*.
- Rudin, C.; and Wagstaff, K. L. 2014. Machine learning for science and society. *Machine Learning*, 95(1): 1–9.
- Salganik, M. J.; and Heckathorn, D. D. 2004. Sampling and estimation in hidden populations using respondent-driven sampling. *Sociological methodology*, 34(1): 193–240.
- Sambasivan, N.; Arnesen, E.; Hutchinson, B.; Doshi, T.; and Prabhakaran, V. 2021. Re-Imagining Algorithmic Fairness in India and Beyond. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’21, 315–328. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383097.
- Santurkar, S.; Durmus, E.; Ladhak, F.; Lee, C.; Liang, P.; and Hashimoto, T. 2023. Whose opinions do language models reflect? *arXiv preprint arXiv:2303.17548*.
- Sap, M.; Card, D.; Gabriel, S.; Choi, Y.; and Smith, N. A. 2019. The risk of racial bias in hate speech detection. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, 1668–1678.
- Selbst, A. D.; boyd, d.; Friedler, S. A.; Venkatasubramanian, S.; and Vertesi, J. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the conference on fairness, accountability, and transparency*, 59–68.
- Shanmugam, D.; Diaz, F.; Shabani, S.; Finck, M.; and Biega, A. 2022. Learning to Limit Data Collection via Scaling Laws: A Computational Interpretation for the Legal Principle of Data Minimization. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’22, 839–849. New York, NY, USA: Association for Computing Machinery. ISBN 9781450393522.
- Shapira, O.; Pasunuru, R.; Ronen, H.; Bansal, M.; Amsterdam, Y.; and Dagan, I. 2021. Extending Multi-Document Summarization Evaluation to the Interactive Setting. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 657–677. Online: Association for Computational Linguistics.
- Sharma, V.; Kumar, N.; and Nardi, B. 2023. Post-growth Human–Computer Interaction. *ACM Transactions on Computer-Human Interaction*, 31(1): 1–37.
- Shirley, K. E.; and Gelman, A. 2015. Hierarchical models for estimating state and demographic trends in US death penalty public opinion. *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 1–28.
- Skalse, J. M. V.; and Abate, A. 2022. The Reward Hypothesis is False. In *NeurIPS ML Safety Workshop*.
- Skalse, J. M. V.; Howe, N. H. R.; Krashennnikov, D.; and Krueger, D. 2022. Defining and Characterizing Reward Gaming. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.
- Sloane, M.; Moss, E.; Awomolo, O.; and Forlano, L. 2020. Participation is not a design fix for machine learning. *arXiv preprint arXiv:2007.02423*.
- Solaiman, I.; and Dennison, C. 2021. Process for adapting language models to society (palms) with values-targeted datasets. *Advances in Neural Information Processing Systems*, 34: 5861–5873.
- Sorensen, T.; Jiang, L.; Hwang, J.; Levine, S.; Pyatkin, V.; West, P.; Dziri, N.; Lu, X.; Rao, K.; Bhagavatula, C.; et al. 2023. Value Kaleidoscope: Engaging AI with pluralistic human values, rights, and duties. *arXiv preprint arXiv:2309.00779*.
- Stanojević, M.; and Sima'an, K. 2014. BEER: BETter Evaluation as Ranking. In Bojar, O.; Buck, C.; Federmann, C.; Haddow, B.; Koehn, P.; Monz, C.; Post, M.; and Specia, L., eds., *Proceedings of the Ninth Workshop on Statistical Machine Translation*, 414–419. Baltimore, Maryland, USA: Association for Computational Linguistics.
- Strathern, M. 1997. ‘Improving ratings’: audit in the British University system. *European review*, 5(3): 305–321.
- Subramonian, A.; Yuan, X.; Daumé III, H.; and Blodgett, S. L. 2023. It Takes Two to Tango: Navigating Conceptualizations of NLP Tasks and Measurements of Performance. *arXiv preprint arXiv:2305.09022*.
- Suresh, H.; and Gutttag, J. 2021. A framework for understanding sources of harm throughout the machine learning life cycle. In *Equity and access in algorithms, mechanisms, and optimization*, 1–9.
- Suresh, H.; Movva, R.; Dogan, A. L.; Bhargava, R.; Cruxen, I.; Cuba, Á. M.; Taurino, G.; So, W.; and D’Ignazio, C. 2022. Towards Intersectional Feminist and Participatory ML: A Case Study in Supporting Femicide Counterdata Collection. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, 667–678.
- Thiago, D. O.; Marcelo, A. D.; and Gomes, A. 2021. Fighting hate speech, silencing drag queens? artificial intelligence in content moderation and risks to lgbtq voices online. *Sexuality & culture*, 25(2): 700–732.
- Thomas, R. L.; and Uminsky, D. 2022. Reliance on metrics is a fundamental challenge for AI. *Patterns*, 3(5).
- Tsing, A. L. 2012. On NonscalabilityThe Living World Is Not Amenable to Precision-Nested Scales. *Common knowledge*, 18(3): 505–524.
- Tufekci, Z. 2014. Big questions for social media big data: Representativeness, validity and other methodological pitfalls. In *Proceedings of the international AAAI conference on web and social media*, volume 8, 505–514.
- Van der Wal, Z.; and Van Hout, E. T. J. 2009. Is public value pluralism paramount? The intrinsic multiplicity and hybridity of public values. *Intl Journal of Public Administration*, 32(3-4): 220–231.
- Varshney, K. R. 2023. Decolonial AI Alignment: Vi\{s\} esadharma, Argument, and Artistic Expression. *arXiv preprint arXiv:2309.05030*.

- Voorhees, E. M.; Craswell, N.; and Lin, J. 2022. Too Many Relevant: Whither Cranfield Test Collections? In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '22, 2970–2980. New York, NY, USA: Association for Computing Machinery.
- Wagner, C.; Strohmaier, M.; Olteanu, A.; Kıcıman, E.; Contractor, N.; and Eliassi-Rad, T. 2021. Measuring algorithmically infused societies. *Nature*, 595(7866): 197–204.
- Wang, Q.; Madaio, M.; Kane, S.; Kapania, S.; Terry, M.; and Wilcox, L. 2023. Designing responsible AI: Adaptations of UX practice to meet responsible AI challenges. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–16.
- Watkins, E. A.; McKenna, M.; and Chen, J. 2022. The four-fifths rule is not disparate impact: a woeful tale of epistemic trespassing in algorithmic fairness. *arXiv preprint arXiv:2202.09519*.
- Wei, J.; Tay, Y.; and Le, Q. V. 2022. Inverse scaling can become U-shaped.
- Wick, J.; Campbell, D. J.; McAlister, F. A.; Manns, B. J.; Tonelli, M.; Beall, R. F.; Hemmelgarn, B. R.; Stewart, A.; and Ronksley, P. E. 2022. Identifying subgroups of adult high-cost health care users: a retrospective analysis. *Canadian Medical Association Open Access Journal*, 10(2): E390–E399.
- Wilkinson, D.; and Thelwall, M. 2010. Social network site changes over time: The case of MySpace. *Journal of the American Society for Information Science and Technology*, 61(11): 2311–2323.
- Winner, L. 1987. *The whale and the reactor: A search for limits in an age of high technology*. University of Chicago Press.
- Wu, S.; and Dredze, M. 2020. Are All Languages Created Equal in Multilingual BERT? In Gella, S.; Welbl, J.; Rei, M.; Petroni, F.; Lewis, P.; Strubell, E.; Seo, M.; and Hajishirzi, H., eds., *Proceedings of the 5th Workshop on Representation Learning for NLP*, 120–130. Online: Association for Computational Linguistics.
- Yoo, Y. 2015. It is not about size: a further thought on big data. *Journal of Information Technology*, 30(1): 63–65.
- Young, M.; Ehsan, U.; Singh, R.; Tafesse, E.; Gilman, M.; Harrington, C.; and Metcalf, J. 2024. Participation versus scale: Tensions in the practical demands on participatory AI. *First Monday*.
- Zelizer, V. A. 1996. Payments and social ties. In *Sociological forum*, volume 11, 481–495. Springer.
- Zheng, A. 2015. *Evaluating Machine Learning Models*. O'Reilly Media, Inc.
- Zhou, P.; Zou, Y.; Yuan, X.-T.; Feng, J.; Xiong, C.; and Hoi, S. 2021. Task similarity aware meta learning: theory-inspired improvement on MAML. In de Campos, C.; and Maathuis, M. H., eds., *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, volume 161 of *Proceedings of Machine Learning Research*, 23–33. PMLR.