Towards Accountability for Machine Learning Datasets: Practices from Software Engineering and Infrastructure

Ben Hutchinson, Andrew Smart, Alex Hanna, Emily Denton, Christina Greer, Oddur Kjartansson, Parker Barnes, Margaret Mitchell

{benhutch, andrewsmart, alexhanna, dentone, ckuhn, oddur, parkerbarnes, mmitchellai}@google.com

ABSTRACT

Datasets that power machine learning are often used, shared, and re-used with little visibility into the processes of deliberation that led to their creation. As artificial intelligence systems are increasingly used in high-stakes tasks, system development and deployment practices must be adapted to address the very real consequences of how model development data is constructed and used in practice. This includes greater transparency about data, and accountability for decisions made when developing it. In this paper, we introduce a rigorous framework for dataset development transparency that supports decision-making and accountability. The framework uses the cyclical, infrastructural and engineering nature of dataset development to draw on best practices from the software development lifecycle. Each stage of the data development lifecycle yields documents that facilitate improved communication and decision-making, as well as drawing attention to the value and necessity of careful data work. The proposed framework makes visible the often overlooked work and decisions that go into dataset creation, a critical step in closing the accountability gap in artificial intelligence and a critical/necessary resource aligned with recent work on auditing processes.

CCS CONCEPTS

• **Information systems** → **Data management systems**; • **Computing methodologies** → *Machine learning*.

KEYWORDS

datasets, requirements engineering, machine learning

ACM Reference Format:

Ben Hutchinson, Andrew Smart, Alex Hanna, Emily Denton, Christina Greer, Oddur Kjartansson, Parker Barnes, Margaret Mitchell. 2021. Towards Accountability for Machine Learning Datasets: Practices from Software Engineering and Infrastructure. In *Conference on Fairness, Accountability, and Transparency (FAccT '21)*, March 3–10, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3442188.3445918>

1 INTRODUCTION

Machine learning faces a crisis in accountability. Deep learning can match or outperform humans on some tasks [94, 101] and is touted

as paving the way for achieving human-level artificial general intelligence [7]. However the datasets which machine learning (ML) critically depends on—and which frequently contribute to errors—are often poorly documented, poorly maintained, lacking in answerability, and have opaque creation processes. This paper argues that the development of ML datasets should embrace engineering best practices around visibility and ownership, as a necessary (but not sufficient) requirement for accountability, and as a prerequisite for mitigating harmful impacts.

Despite rapid growth, the disciplines of data-driven decision making—including ML—have come under sustained criticism in recent years due to their tendency to perpetuate and amplify social inequality [13, 44]. Data is frequently identified as a key source of these failures through its role in “bias-laundering” [40, 51, 54, 119, 125]. For example, recent studies have uncovered widespread prevalence of undesirable biases in ML datasets, such as the underrepresentation of minoritized groups [27, 40, 131] and stereotype aligned correlations [28, 51, 72, 155]. Datasets also frequently reflect historical patterns of social injustices, which can subsequently be reproduced by ML systems built from the data. For example, in a recent study examining the datasets underlying predictive policing models deployed in police precincts across the US, the underlying data source was found to reflect racially discriminatory and corrupt policing practices [119]. The norms and standards of data collection within ML have themselves been subject to critique, with scholars identifying insufficient documentation and transparency regarding processes of dataset construction [52, 53, 126], as well as problematic consent practices [114]. The lack of accountability to datafied and surveilled populations as well as groups impacted by data-driven decisions [32] has been further critiqued.

Taken together, these objections around data raise a serious challenge to the justifiability of the datasets used by many ML applications. How can AI systems be trusted when the processes that generate the their development data are so poorly understood? If ML is to survive this crisis in a responsible fashion, it must adopt visibility practices that enable accountability and responsibility throughout the data development lifecycle.

This paper takes a step in this direction. We argue for a rigorous and visible dataset creation process that is based in foundational ontological thinking about what kind of “things” datasets are. We posit that datasets constitute a form of *technical infrastructure*, in the sense that they are enabling and necessary components of technical systems, with which they interact through technical interfaces. Datasets are also *infrastructure* in the infrastructure studies sense, in that they arrange a way to view and structure knowledge of the world [22, 88]. We show how framing datasets as infrastructure helps to shed light on many of the social and technical phenomena

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

FAccT '21, March 3–10, 2021, Virtual Event, Canada

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8309-7/21/03.

<https://doi.org/10.1145/3442188.3445918>

that we see around data; for example, datasets are critical components of ML development but their existence and the conditions of their creation are often invisible and hence undervalued. Further, given their nature as technical infrastructure, the processes that give birth to datasets are best understood as fundamentally goal-driven engineering. We argue that debates around ML data are sometimes complicated by confusions around the epistemological goals of engineering. This is in part due to computer science’s own methodological complexity, using varying techniques of knowledge validation that include formal proof, automated testing of both deterministic and stochastic forms, user testing, as well as techniques aimed at providing explainability, traceability and interpretability.

This paper explores in detail a number of the practices that need to be adopted to mitigate dataset risks. These practices in turn require adopting a deliberative and intentional methodology, rather than the *post hoc* justifications that are sometimes observed when datasets are developed hastily or opportunistically. From an observer’s perspective, these practices replace *abductive* and *a posteriori* reasoning about dataset provenance with careful documentation, facilitating reviews and audits [115]. They help us to avoid the trap of being blindly “data-driven” by poorly-understood datasets—mistaking taking a back seat with “empirical rigour”—when it is in fact humans who we need to be clearly in the driver’s seat in ML development, setting goals, direction and strategy with deliberation and intentionality. These practices thus provide bounds on accountability both in dataset creation, and also of the ML systems that depend on those datasets as infrastructure.

We stress the importance of establishing a **non-linear cycle of dataset development**, with analysis, design and evaluation central to the process (see Figure 1). Building off this lifecycle, key practices which we will cover include:

- **Documentation:** a model of **documentation practices** throughout the dataset development lifecycle, drawing on software lifecycle practices (see Table 1 and Section 4.2);
- **Oversight:** **diverse oversight processes**, including audits and reviews, which leverage these practices (see Section 4.3);
- **Maintenance:** **robust maintenance mechanisms**, including those for addressing technical debt, correcting errors, postmortems and adapting to changing contexts (see Sections 4.2.5 and 5.1).

The structure of this paper is as follows. We first make the case for datasets being a form of technical infrastructure (Section 2). Then, by considering the epistemologies, cultures and methods of engineering, we argue that dataset development is essentially a form of engineering practice (Section 3). Building on this, we show how software engineering and infrastructure practices provide lessons for dataset accountability (Section 4), including a model for dataset development documentation. We then discuss the implications of this model for broader systems and ecologies in which datasets are maintained and used, proposing research directions and argue for cultural shifts (Section 5).

2 DATASETS AS INFRASTRUCTURE

ML fairness practitioners have identified collection and curation of representative data as the most common way to mitigate model

Figure 1: The Dataset Development Lifecycle requires documentation for each stage. See Table 3 for descriptions of each stage, and Table 1 for document types.

biases [71], echoing the broader tendency within ML communities to valorize larger datasets as necessarily better—exemplified by [62] and subsequently in [85, 138]. However, datasets are not merely a collection of “facts” to be discovered, but rather construct a particular kind of knowledge (see for example [90]). Given the centrality of data to the field of ML and its criticality to the fairness of ML-based systems, it is useful to consider two questions, one ontological and one explanatory: *what kinds of things are datasets?* and *why do datasets so often seem to have problems?*

The answers to both these questions are intertwined, and are inseparable from the processes of dataset creation and use. Being essential to ML, datasets have all the classic properties of IT infrastructure: they are shared information delivery bases with well-defined architectures and standardized interfaces, creating value by enabling rapid iteration and new forms of processes, from which their ultimate value is derived [19, 122]. The dependency is in fact so close that in common parlance the ML model and its training data are sometimes used interchangeably, for example biases in models are frequently seemingly identified with training data bias, or used to abductively infer biases in the training data (e.g., [11, 38, 77]).

Despite the critical role data plays, the systemic devaluation of data work is in plain view. Enormous conferences, prizes, and adulation await new algorithmic achievements on “benchmark” datasets (seen as fixed resources, or *data* in the original Latin sense [50]). In contrast, datasets, like other infrastructures, are often taken for granted after they are built, and fade into the background until they break down [22, 41]. Often playing the role of commons [49], benchmark datasets exhibit “comedy of the commons” whereby usage of a dataset increases its value, due to the comparisons enabled by its role as a measuring instrument [146], while the patient and careful work of caring for datasets (including the collection, cleaning, and annotation of data) receives few, if any, accolades in the mainstream field. Since it is often an unpaid burden on top of existing workloads [73, 125, 140], careful dataset work is unlikely to occur at all. Publications focusing solely on datasets rarely make it through

DATASET ACCOUNT	QUESTIONS ANSWERED BY THE DOCUMENT	KEY ROLES
Dataset Requirements Specification	Is data needed? By whom? What are its intended uses? What properties should it have? Do the uses necessitate constraints on collection and/or annotation?	Requirements owner; stakeholder; reviewer
Dataset Design Document	Is a new dataset needed (do existing datasets not meet requirements)? How will requirements be operationalized? What tradeoffs & assumptions are being made?	Design owner; domain expert; reviewer
Dataset Implementation Diary	How were the design decisions implemented? Why were they done this way? What unforeseen circumstances arose when implementing the design?	Implementation owner; data creator/labeler
Dataset Testing Report	Should I use the dataset? Does the dataset meets its requirements? Is the dataset safe to use? Is the dataset likely to have previously unforeseen consequences?	Data scientist; adversarial tester
Dataset Maintenance Plan	How will data staleness be detected and fixed? How will errors be fixed? How will affordances for manual interventions be provided?	Maintainer; funder; bug filer; data contestor

Table 1: Critical document types for accountable dataset development. Each one is directly analogous to documentation types produced by the Software Development Lifecycle.

the review process of top-tier venues [67] and—perhaps due to the imperative to squeeze dataset details into methods-oriented papers—critical dataset design decisions are frequently left underspecified (e.g. [53]). Despite the foundational role data place in ML research and development, data occupies little to no space within predominant ML textbooks and data is not recognized as an area of ML specialization [75]. ML data are also not version-controlled, held at institutional repositories, or given DOIs; and data annotation work is given short shrift in discourses about model building and development.

Thus, it is not solely the invisibility of the dataset to end users (for the learning algorithm is also invisible) that alone explains the problems we see in datasets, but also the devaluing of the processes of dataset creation. And here a vicious circle enters the picture. Because benchmark datasets shift the field’s attention to differences between learning algorithms, as measured by narrow benchmark metrics, we see an emphasis on explaining and documenting those differences. In contrast, the benchmark dataset, now a fixed resource for measurement and competition, has little need for explanation or documentation. For datasets, poor valuation leads to poor documentation, which in turn leads to even poorer valuation as the “fixed” becomes the “taken for granted”. An ML dataset, after its release, quickly becomes a “black box” of scientific inquiry, but possibly without facing requisite “trials of strength” that technoscience tools in other domains may encounter [89].

In many respects, then, dataset development shares inherent similarities with work on other forms of technical infrastructure. It is deep in the training of computer programmers to take infrastructure for granted; indeed, the use of abstractions in software libraries encourages this (the “portability trap” of [129]). Although software developers and model developers both need to “get things done” by working with abstractions, they also need processes of justifying and/or fixing what’s behind the infrastructural curtains, lest abstractions launder unjustified assumptions. We see that work on computing technical infrastructure faces similar challenges with respect to valuation and appreciation. While product software is directly identified with the “product”, the networks, datacenters, release management tools, and testing frameworks that are critical to deployment remain (when successful) invisible and taken for granted, or (when not) blamed for failures. Corporate value is assigned to the product software—and measured through clicks,

views, *etc.*—while improvements to shared and distributed technical infrastructure are much harder to assign value to.

This extended comparison with infrastructure highlights patterns of valuation that are operative in both ML and engineering domains. By considering the ways in which practices of software infrastructure have fought against devaluation, we see opportunities to increase the valuation of dataset labor. Improvements on many fronts are needed, including changing the practices of academia to view data as the first-class infrastructure of ML: data should be critical to the ML curriculum; papers about data should be recognized and rewarded in ML fora; data should be interrogated rather than taken for granted. A further, but critical, change involves reversing the vicious circle described above, and through robust dataset documentation practices bring more attention to the processes and value of dataset labor. In Section 4.2 we discuss in detail such documentation processes, but first we explore the dataset development process in more detail.

3 DATASET DEVELOPMENT AS ENGINEERING

The invisibility of dataset processes described in the prior section leads to a lack of critical examination of datasets used to build models. Are researchers and engineers justified in trusting the datasets that they use to build their models? Are datasets fit for their intended purposes? Do the datasets contain hidden hazards that can make models biased or discriminatory? What are the human biases that have been built into the datasets? (And what are the implications if your model performs well on a biased dataset?) Understanding the processes by which datasets are conceived and constructed is central to these questions. Specifically, in the systems of dataset creation, what are the processes, how are these processes documented, and *who is responsible for what?* Even more fundamentally, what are the cultures of ML data work that enable these systems? In this section we make the case that dataset development systems are essentially engineering systems, in their goals, practices, solutions and cultures. By considering the nature of engineering, we identify five arguments that dataset development is foundationally a type of engineering.

3.1 Datasets provide “knowledge-how”

What distinguishes engineering from other technical and empirical disciplines, and especially from scientific disciplines typified by the physical sciences? Despite widespread assumptions to the contrary, engineering is not simply applied science, for “while science aims for knowledge, engineering aims for useful change” [26]. The distinction is knowing how to accomplish specified goals, in contrast to knowing that something is true; in other words, while science aims for *knowledge-that*, engineering aims for *knowledge-how* [123]. As such, science is typically concerned with explanation, certainty, universality, abstractness and theory; while engineering is characterized by contingency, probability, particularity and concreteness [57]. The two are interdependent: engineering can be in support of science (e.g., datasets can be used towards scientific goals); and modern science is heavily dependent on engineering [24], and on big data [92]. However, due to their infrastructural roles (see Section 2), the development of datasets is fundamentally a form of engineering. Datasets are designed for specific purposes, and their usefulness must be judged with regard to that purpose [150]. As with engineering, good practice when developing datasets demands careful attention to relevant science (including social sciences and engineering sciences). Indeed, the consultation of relevant science is a cornerstone of ideal engineering.

3.2 Dataset practices are political

Given that engineering knowledge is *knowledge-how*, it is critical to dispel the myth (discussed in [59]) that engineering practice is a “purely technical” matter. Engineering artefacts are imbued with politics via their roles in sociopolitical systems [57, 149], and so too are ML datasets inseparable from politics [36, 41, 114]. To be “just an engineer” is to acknowledge one’s role, stakes, and domains of expertise within a sociotechnical system, while also acknowledging the roles, stakes, and expertise of others. It requires accountability for enabling objectives, even if those objectives are not one’s own. Within information science disciplines, engineering roles are filled by data experts, and Ben Green has suggested a four-stage path for how data experts can meaningfully engage with society [59]: becoming interested in addressing social issues, recognizing the underlying politics, redirecting existing methods, and developing practices and methods for working in partnership with communities. Critically, the second stage requires that skills in critical reflection are as essential to dataset development as knowledge of statistics. This aligns with the prescriptions of Neff *et al.*, who ask data scientists to consider several interpretative and value-laden critiques of their data while working with it [105].

3.3 Datasets are models of reality

Central to engineering is the concept of *engineering models*, which are representations of reality used for design and testing prior to real-world deployment [26]. Engineering models are abstracted simplifications of reality, necessarily incomplete, emphasizing some features of reality while deemphasizing or overlooking others [145]; they are plausible aids but “potentially fallible” [82]. ML datasets are in fact prototypical engineering models, representations of facts about the world that cannot be experienced directly, nor often replicated. Like other engineering models, datasets can be developed

with a number of different end goals including explanation, development of intuition, instruction, prediction, design, evaluation and experimentation [3], and their goals dictate priorities and tradeoffs between capturing different dimensions of reality. Uncertainty is a fundamental feature of all models, including datasets [26], and robust testing and calibration is required to avoid brittle and unsafe conclusions. Bulleit *et al.* report that in engineering disciplines “designers rarely get feedback on how well their models represent the specific full-scale projects that they have designed” [26], and the infrastructural properties of data discussed above suggest the same is true of ML dataset workers.

3.4 Dataset development is non-linear

Engineering problems “tend to be ill-structured” [134] or even “wicked” [120], meaning they defy deductive reasoning and definitive answers. Instead, engineering involves reconceptualizing a complex situation to facilitate analysis; with feedback loops impacting problem (re-)definition (see [109]), rather than focusing narrowly on problem solution. As goal-oriented models of reality, dataset development requires iteration and feedback loops due to changes in stakeholders and goals, as design assumptions become invalidated or better alternatives become available, and fundamentally as reality itself changes. Western culture has devalued engineering with respect to science (albeit not social science) precisely because of engineering’s context sensitivity, pluralism, and temporal instability [57], mirroring data work’s devaluation with respect to ML model work (Section 2).

3.5 Cultures of dataset work mirror those of software

Computer programming emerged from an intersection of multiple academic disciplines, and the intrinsic cultures of these fields informed the norms and attitudes that persist in software engineering today [111]. The rise of ML and its big, shared datasets are subject to some of the same cultural pressures and fault lines; their solutions to collective coordination and communication challenges giving rise to similar cultural norms. Indeed, the four cultures of programming identified by [111] are evident when looking at the development of ML models and datasets:

Mathematical culture: This is evident in the focus on mathematical and statistical methods at the heart of ML, the value given to formal proofs, and the tendency of practitioners to focus on tractable data problems such as agreement and reliability, rather than more complicated data questions such as external validity (e.g., whether arrests represent crimes) [74, 146].

Hacker culture: This culture is expressed in ML via a focus on individual contributions by “superstars” [60], “savants” [117] and saviours with brilliant innovative insights. Hacker culture emphasizes moving fast and debugging as errors arise [91] (neural network interpretability methods [31] are the stack trace debuggers of ML). This culture devalues the incremental and cooperative care required to create high quality datasets, and has been implicated in the discriminatory effects of technological systems.

Managerial culture: Appreciating software development as inherently sociotechnical, this culture focuses on team structures emphasizing responsibility, shared documentation and auditable

BARRIER	DATASET CONCERNS	PROPOSALS FOR MITIGATING
Problem of many hands	Datasets teams may have indeterminate structure and responsibilities; data may be sourced from third parties, or “from the wild”; data labeling teams may be unknown	Factor dataset development into discrete stages; assign clear responsibilities to the owner of each stage; robust documentation serving as boundary objects
Unintended artifacts	Intended properties of dataset may not be explicit; datasets may contain unforeseen problems	Consult diverse stakeholders; document intentions explicitly; assume data is “guilty until proven innocent”; adversarial testing
Computer as scapegoat	“Data as scapegoat”; sourcing data from logs of applications; sourcing data from generative models	Require visibility into conditions of dataset creation; document and justify decisions to use data from system logs; understand and document properties of generative models

Table 2: Three of Nissenbaum’s barriers to accountability [106], specific data concerns, and proposals in this paper.

processes [52, 99, 115], and advancing these practices for ML data is a core goal of this paper.

“Engineering” culture: Dominated by practical concerns, and accepting the impossibility of formal guarantees, engineering cultures of software development focuses on reliable tools and practical methods. Confidence in systems is gained not by proof or debugging, but by error measurement and error handling. In ML data, the emphasis on shared benchmark datasets is a cornerstone of ML engineering culture [146], and “Humans in the Loop” in ML deployments (as a form of robust error handling) also fall under this culture.

These correspondences between the cultures, goals and methods of datasets and engineering point the way forwards. In the following section we explore how best practices in software development can be adapted to improving the visibility and quality of dataset development.

4 TOWARDS ACCOUNTABILITY FOR DATASETS

Accountability has been described as “fundamentally about answerability of actors for outcomes” [83]. Critical to this question is what information is known by whom, and how it is used. In fact, the sharing of information is the first key phase of mechanisms of accountability [147]. The second phase involves deliberation and discussion by the “forum”, as well as requests for more information. During the third and final phase, the forum imposes consequences. In this section we focus our attention on the first, critical, information phase, and ask two questions: 1) *which actors* should provide information about datasets? and 2) *what information* should these actors provide to enable meaningful discussion and consequences?

The information that is shared as a necessary (but not sufficient) precondition for accountability is referred to technically as *accounts*. The recording of dataset accounts is at its most fundamental a question of bookkeeping, but the details are critical: which books should be kept, what are their stories, and who are their authors? The answers lie in the closely related goals of traceability, attributability, responsibility, accountability, and answerability, the distinctions between which have been the subject of much discussion in Western philosophy. Judgements of fault can be on multiple dimensions of responsibility, e.g., under one account: “attributability-responsibility” concerns to an agent’s character, “accountability-responsibility” refers to an agent’s regard for others,

and “answerability-responsibility” is about an agent’s evaluative judgments [132, 133]. In this paper, we have little to say about judgments of character, other than calling for transparency about which actors bear attributability-responsibility. (Reliability is clearly a concern here: does this institution/individual have a history of developing trustworthy and responsible artefacts, cf. [135]?) However, below we will detail how both regard for others and evaluative decisions are things for which accounts can and should be kept when developing datasets.

Helen Nissenbaum [106] describes barriers to accountability of computer systems which are all directly applicable to datasets. Table 2 summarizes how her concerns relate to datasets, as well as the mitigating proposals on which we will elaborate below.¹ Our analysis and proposals follow [147] and [83] in using temporal situatedness as a key organizing factor. At the highest level, what are the dataset accounts that should be kept before (*ex ante*), during (*in medias res*) and after (*ex post*) the collection and labeling of data? Based on the robust parallels with software engineering argued for in Section 3, we propose that the software development lifecycle provides an appropriate model for the dataset development lifecycle. This builds on previous analyses of the machine learning lifecycle (e.g., [9, 153]) and the data science lifecycle [21], by exploding what is sometimes considered a mere single stage of “Data Management”.²

4.1 Lessons from engineering models

Alvi proposes a set of recommendations for engineering models (Section 3) [3], and they are equally applicable to datasets. We summarize them here, substituting “dataset” for “model” for clarity.

- Treat datasets as *guilty until proven innocent*.
- Identify the assumptions underlying a dataset in writing.
- Conduct independent peer reviews and checks during and after dataset development.
- Evaluate datasets against experience and judgment.
- Use visualization tools.
- Perform sensitivity studies of dataset parameters.
- Understand the assumptions and limitations of datasets, rather than using them as black boxes.

¹Nissenbaum’s fourth concern is about liability, which is beyond the scope of this paper.

²We differ from the “data lifecycle” analysis of [113] in that they consider the “first step of ML is to understand your data”, presuming that data already exists.

<i>Requirements analysis</i>	Deliberations about intentions, consultations with stakeholders, and analysis of use cases determine what data is required.
<i>Design</i>	Research is performed and subject matter experts are consulted in order to determine whether the data requirements can be met, and if so how best to do so.
<i>Implementation</i>	Design decisions are transformed into technologies such as software systems, annotator guidelines, and labeling platforms. Actions may employ and managing teams of human expert raters.
<i>Testing</i>	Data is evaluated and decisions about whether or not to use it are made.
<i>Maintenance</i>	Once collected, a dataset requires a large set of affordances, including tools, policies and designated owners.

Table 3: Stages of the Data Development Lifecycle.

The model of dataset development documentation that we present in the following section is aligned with these recommendations.

4.2 Documentation of Dataset Development

Documentation of the stages of the dataset development process serve as *accounts*, in the sense just described above. As such, we now delve into the documentation artefacts relevant to each of the lifecycle stages (Table 1). Although the precise number and nature of lifecycle stages is sometimes debated, for the current purposes it is sufficient to consider the 5 stages of the lifecycle in Table 3. Each document has owners with designated roles and duties [84].

4.2.1 Data requirements accounts. Why is data the solution to a problem at hand? Given that data is infrastructure (see Section 2), and that datasets are engineering artefacts (Section 3), it is important to consider the ends to which data is the means, and (to avoid data-solutionism) whether data is the best way of addressing this problem. Given the long history of existing data being put to novel uses, one must also account for the consequences of unintended uses. The thoroughness of such accounts speaks directly to questions of *ex ante* accountability. The 6 W's provide a useful framework for detailing such intended and unintended scenarios, *i.e.*, specifying for each use the who (the decision-makers) and the whom (those being impacted), as well as the what, where, when and why [147].

In engineering, these considerations all fit under requirements analysis [12], and we embrace the proposals of [96, 97, 142] for **Data Requirements Specifications**—directly analogous to software requirements—covering both quantitative and qualitative factors. One of the key tasks is holding stakeholder requirements development sessions in which the (possibly conflicting) needs of multiple stakeholders are collected and aggregated, and conflicts resolved through accountable mechanisms (including the keeping of accounts regarding the conflicts), making transparent what is

valued in the data [143]. As with software requirements, a number of challenges may exist when documenting dataset requirements. Stakeholders may not communicate their needs, or may be reluctant to commit to a set of written data requirements. Data development teams may have a natural inclination to “jump in” and start collecting data before the requirements analysis is complete. Stakeholder groups may have different vocabularies, or may be unable to reach consensus using previously agreed decision-making mechanisms. Stakeholders may try to make the requirements fit an existing dataset, rather than think through their specific needs. Cognitive biases and standpoints may prevent complete accounts.

The task of producing requirements documents should have a clearly designated owner. The frequent neglect of clearly specified requirements for data goes beyond ML data (*e.g.*, [33]). There is likely a need to encourage and reward the development and ownership of requirements specifications, by making this part of a key task, rather than a new burden placed on top of already demanding schedules [125].³ Since the benefits and risks of infrastructure accrue over time, the long horizons of the impacts of data requirements analysis need to be taken into account when judging success or failure. In other words, both the quality and quantity of uses of the dataset, over long time horizons, need to be taken into account.

Appendix A provides a flexible template for documenting dataset requirements, adaptable to many contexts and organizational environments using processes of contextual inquiry [15]. It covers a broad range of requirements about data instances, distributions, processing and sourcing, motivation and intended uses, as well as critical metadata such as approvers and changelog.

4.2.2 Dataset design accounts. The distinct roles of requirements and design are sometimes misunderstood,⁴ but it is important to emphasise to stakeholders the conceptual distinction between requirements and design, to avoid premature decisions about design before requirements are established. Whereas requirements analysis sets the goalposts for the dataset, the design answers the question of means: *how will we get there?* That is, the former answers *what?* and *why?* by documenting ends and intentions, while the latter answers *how?* by documenting means. To give a software analogy, while a requirement may be that a function return a sorted list, a design decision may be to use the bubble sort algorithm. Research of existing options is a critical part of the design phase. If a pre-existing dataset (or combination of datasets) is satisfactory, there is usually no good reason to design and build a new one. This stage may therefore necessarily involve elements of Dataset Testing (see below), as pre-existing datasets are evaluated against the current requirements.

The primary account of this stage is the **Dataset Design Document**. This document's primary roles are to lay out the plan of how requirements will be achieved, and to justify the design decisions that are made. These justifications take myriad forms, but consultation of domain experts is a critical part of this stage. Crystal clear and objective options are often lacking, but rather the formation of information into discrete cases involves decisions

³This burden is common to many fields: “faced with ever more crowded diaries, we increasingly subsidize data collection to a very significant degree through our personal time and often through enormous personal sacrifice” [87].

⁴And as [86] drily observe, machine learning is characterized by a lack of both requirements specifications and design specifications.

about standardization and lossiness [29]. Since the future lives of datasets are impossible to predict (and “researchers in neighboring fields may view the same data in considerably different ways” [29], see also [108]), creating detailed accounts of assumptions and decisions is critical to the dataset design phase. Documenting trade-offs and which alternatives were considered—both technical and nontechnical—are central to the design process, and integral to engineering practice [58]. Even seeming “neutral decisions”, such as how to sample, require multiple decisions: datasets require formation of cases, which require categories, and these categories have consequences that are frequently statistical and sometimes moral. Aspects of phenomena which are easily calculated become amplified, while those which resist standardization are reduced or excluded cited by [29].

The field of Natural Language Processing’s historical use of carefully constructed balanced datasets—prior to the proliferation of internet data [79]—points to lessons about dataset design which remain relevant. These include the necessity of detailed discussions of “what it means to ‘represent’ a [phenomenon]”, “definition of the target population,” and consideration that “theoretical research should be prior” in dataset design “to identify the situational parameters,” after which construction of the dataset “would then proceed in cycles” [20].

Appendix B provides a flexible template for documenting dataset design decisions, adaptable to many contexts and organizational environments. It includes sections detailing related datasets, how data will be sourced and annotated, key empirical characteristics, how privacy will be handled, as well as how data quality will be measured.

4.2.3 Dataset implementation accounts. Regardless of the combination of human labour and computation involved in dataset creation, more decisions have to be made—and documented—during dataset implementation. To extend the previous comparison, if the design phase specifies the bubble sort algorithm is to be used, the implementation might choose data structures and control flows whose appropriateness may be contextually dependent. An analogy of dataset implementation accounts is therefore with code comments: fine-grained documentations tightly coupled with the implementation. Since digital traces can often be mined for the answers to *what was done*, as with code comments the most important implementation accounts are for explaining *why things were done that way* [139]. There are many ways that these accounts of data collection and labeling can be recorded. As with code comments, tight coupling has benefits of synchronicity; that is, tools that are used to manage and support data collection and labeling should ideally have integrated commenting capabilities, with comments being exportable and searchable. Failing such tight integration, a temporally organized **Dataset Implementation Diary**, can be used. Issue-tracking systems can also be adapted for this role, e.g. with each decision being blocked until justified on the appropriate “ticket”.

4.2.4 Dataset testing accounts. Data needs testing just like code does [142]. As with software testing, there are many different flavours of dataset testing, but at their foundation is the question *should this data be used?* A **Dataset Testing Report** is the account that specifies the evaluations that were done, as well as their results.

These are naturally analogous to software unit testing and security testing. In both cases, accounting of the diligence and coverage of the test cases provides the “receipts” for justifying trust in the data. At the lower levels of testing are questions of data wellformedness (for which a range of tools are available [43]); at higher levels, questions of fidelity [18]. Two particularly important types of testing are requirements testing and adversarial testing.

Requirements testing (also called “acceptance testing” [56]) checks a dataset against the stated requirements for its use. (These can be the requirements which led to the creation of the dataset, but alternatively a pre-existing dataset may be evaluated against a novel set of requirements.) Requirements testing directly reframes the question of “how to evaluate whether the right data was collected with sufficient quantity” [121], by checking whether we got what was needed. As with software testing, ideal data testing returns boolean pass/fail outcomes; difficulties in doing so may point to a lack of clarity in the dataset requirements. However it is not necessary that all inputs to testing be present in the data records themselves; evidence can also come from metadata generated throughout the data collection lifecycle. For example, an intention to share a new dataset may lead to certain requirements in how consent is obtained, and the data collection process needs to keep “the receipts” as evidence of due process. To facilitate this, it is therefore necessary that dataset requirements be *traceable* [100], and in many cases requirements should be measurable. Non-boolean test outcomes are still useful in flagging concerns, and frequently this can lead to discussion and iteration on requirements. When requirements and their testing is specified clearly enough, opportunities for test automation become possible, including realtime monitoring and alerting [23]. When tests have sufficient validity and coverage, automated testing can provide reliable guarantees for continuous data deployments. Automated tests can also feed back into the data collection stage, providing signals as to which parts of the data distribution need further sampling. When such tests involve ML model evaluation, this takes on the flavor of active learning [130].

Adversarial testing of a dataset (also known as “test-to-fail” testing [110]) aims to uncover unforeseen harms which may come from its use. Just as in software security, the tester plays the role of an “attacker” and tries to break things by finding risks of undesirable outcomes. Such risks vary widely in form, from harms to specific individuals such as privacy leaks, to harms to subgroups such as unforeseen correlations or stereotypes, to public relations risks via embarrassing data (including errors or omissions), to the possibility of malicious third parties using the data for nefarious ends.

Appendix C contains a template for Dataset Testing Reports, to summarize the processes and results of testing.

4.2.5 Dataset maintenance accounts. If you can’t afford to maintain a dataset, and you also can’t afford the risks of not maintaining it, then you can’t afford to create it. However a common challenge with many forms of infrastructure is that funding for creation is often more readily available than for maintenance, leading to a maintenance budget gap. Furthermore, charitable volunteerism cannot be expected as a substitute, since maintenance is one of the least desirable roles in the software lifecycle [30] and the same is likely true of data. Indeed since many maintenance tasks result in no

visible impacts beyond preservation of the *status quo*, maintenance work risks being viewed as an unrewarded, unglamorous burden.

A forward-looking **Dataset Maintenance Plan** is the appropriate account for this stage (analogous to a software maintenance plan [2]). Although adding new requirements is sometimes called “perfective maintenance” [17], we consider that task to be a fresh iteration of the requirements phase of the lifecycle; the maintenance plan should instead focus on corrective, adaptive and preventive goals [17, 34]. *Corrective maintenance* aims at fixing errors. Unforeseen problems are to be expected, and affordances for human-data interaction should thus be provided. Data poisoning may have occurred, labeling guidelines may have been misconstrued, or the phenomena in question may quite simply change. Affordances for manual edits and batch updates need to be in place (with appropriate access controls), and all such actions need to be scrupulously logged. Making these logs human-readable and searchable is an important step in dataset maintenance accountability. *Adaptive maintenance* aims at preserving the key properties of a dataset under changing external conditions. The critical question is thus: *what properties should be preserved?* Here requirements come into play: one key goal is preserving the dataset’s satisfaction of its own requirements. Challenges arise when drift occurs, *i.e.*, the dataset diverges from the reality of the phenomena it purports to measure. Addressing this requires feedback loops with other lifecycle stages, to set up recurrent measurement of the validity and fidelity of the data. *Preventive maintenance* aims to prevent problems before they occur, for example paying down technical debt before it becomes unwieldy (see Section 5).

Datasets should be stored or maintained in stable, institutional repositories that allow for differential levels of access and stable universal identifiers. While this may be less of an issue in organizations with mature internal data infrastructures (such as corporate firms), this can present a special problem with datasets which are ostensibly released for use for the research community, and yet are hosted on personal and lab websites. In the research space, stable institutional repositories used for social science research include the Harvard’s Dataverse [80], University of Michigan’s ICPSR [45], and New York University’s Databrary [55]. These are robust projects with years of institutional support, which decreases the infrastructural cost for hosting by lower-resource organizations and teams.

Maintenance costs need to be continually weighed against benefits, and eventually there may come a time when a dataset should no longer be maintained but instead deleted. The easy replicability of binary data means that deleting all copies of a dataset becomes impossible once it has been shared. So in the absence of complex cryptographic solutions, dataset sharing should be tightly controlled whenever contestability is required (see Section 5.1).

4.3 Dataset Audits and Reviews

As discussed above, accountability for datasets to the forum requires discussion and deliberation of the dataset documentation (*i.e.* “accounts”). One type of deliberation process is the audit, which in ML contexts might review datasets and/or models. In internal ML audits, initial scoping and mapping phases are followed by an artefact collection phase, prior to evaluation and reflection by the

audit team during which the artefacts are analysed in light of the organization’s policies and commitments [115]. The dataset development documentation described above form part of the audit trail [48], boosting auditability by design [157]. The audit is but one form of review, however, and just as for software the entire dataset lifecycle requires reviewing (*e.g.*, requirements reviews, design reviews, *etc.*), especially during initial development, but also when contextual circumstances change.

4.4 Lessons from Infrastructure Governance

Governance mechanisms for infrastructure can inform accountability mechanisms for datasets. We focus here on public infrastructure governance, but models of private dataset governance also exist [97]. Like dataset projects, public infrastructure projects have often begun with poor understanding of risks and challenges, and with an underappreciation of the work required. This spirit of entrepreneurialism risks dangerous adventurism [6], and misestimations of both costs and benefits [46]. To mitigate these, risk anticipation and management should be prioritized [6]. One governance mechanism is the stakeholder consultation process in dataset requirements analysis, however a realistic understanding is required of how contested decisions are resolved [144]. One tradeoff is between hierarchical coordination (appeal to institutional authority) and negotiated coordination (*e.g.*, by requiring consensus). The former often involves appeals to some form of welfare-maximization, while the latter may restrict the decision space to Pareto optimal solutions. It may take years for problems in datasets to be discovered (*e.g.*, Imagenet was developed in the late 2000s, with certain issues not identified until 2019 [37]). This supports calls for social accountability mechanisms which emphasise legitimacy via inclusivity in analysis and decision-making phases [76], including focus groups, town meetings, public hearings, and randomised citizen involvement. Identification of “territories” of interests enables conversations about how risks and benefits are accrued to various groups. Temporal discounting and inter-generational effects may also be relevant: are there future effects of datasets on those who are too young to have a say today?

5 DISCUSSION

In Section 4.2 we presented a model of dataset development documentation which facilitates transparency, motivated by accountability’s dependency on visibility. We now discuss the broader systems of dataset development itself, and the cultural systems in which dataset work is embedded.

5.1 Benefits of Data Documentation

Just as with software development [156], there are diverse costs and benefits as to how dataset development documentation interacts with project comprehension, management decision-making, and time efforts/savings. More empirical evaluations of these are needed for data documentation. Improved documentation helps developers with all stages of the lifecycle but, just as with software documentation [156], we expect that improved documentation will particularly aid dataset developers in the maintenance phase, due to both collective amnesias and employee turnover.

Data Work Recognition. Documentation increase employees’ valuation of dataset work by “outlining the nature of an individual

or team’s contribution to an overall system”, creating space for reflection, debate and recognition, and by giving “meaningful understanding of the impact of their personal participation” [116]. Open data advocates have argued that data availability has led to increased data citation, which can serve as a professional incentive for doing necessary data work, as well as developing data documentation [112].

Dataset Productionization. Productionized deployments of ML models require much more than evaluating model performance metrics in isolation [23]. In these contexts, the model itself will form only one part of the deployed system, requiring infrastructure for deployment and monitoring [14]. Dataset development documentation outlined above plays a critical role in defining the contracts between different components, helping product developers, maintainers and end-users determine how much trust to place in the system [8]. Additionally, in online “continuous training, continuous serving” systems, the concept of a “dataset” as a static artifact (or even a versioned artefact) breaks down. Instead, datasets are defined by automated pipelines of data processing steps, requiring automated testing [127] against requirements and expectations, along with monitoring and alerting.

Data Debt Mitigation. The deferred engineering maintenance cost that occurs when speed of execution is prioritized over quality of execution is known as “technical debt”. Like financial debt, technical debt compounds, ultimately slowing down long term progress in pursuit of short term gains (sometimes leading to the “Winner’s Curse” [128]). Technical debt is equally relevant to datasets, in fact dataset dependency debts can be more costly to maintain than code dependencies [127]. Critical to addressing this are maintenance practices that include budgeting for debt repayment, as well as explicit definitions and measurements of quality, via data requirements documents and data testing reports. Undesirable data feedback loops and unutilized data features can be mitigated with careful requirements practices.

Data Postmortems and Premortems. Even the best designed systems occasionally fail [25, 120], however cyclical development practices help avoid making the same failures twice, and the postmortem is a critical tool in learning from past failure and building institutional memory [42]. Institutional debugging of “what went wrong” with datasets requires both a shared definition of what “right” looks like (*i.e.*, requirements), as well as rigorous testing and test validation (for faulty tests are worse than no tests [107]). A premortem involves imagining in advance what could go wrong [81], and, similarly for datasets, requires a shared definition of “right”. Indeed, premortems can be a useful tool during data requirements analysis for identifying gaps and vagueness.

Transparent Dataset Reporting. Transparent reporting of machine learning datasets [16, 52, 70] and the models trained on those data [99] enables third party agents, including potential users, to make informed decisions. Poor reporting practices result in organizations not having the tools they need to understand ML datasets prior to building models [70]. Careful documentation throughout the lifecycle creates the intra-institutional memory required for transparent reporting to third parties.

Data Reuse and Reproducibility. Open reporting of machine learning datasets allows third parties, such as auditors and researchers, to attempt replication of results and predictions, thereby strengthening trust in quantitative findings. This work coalesces well with the movement within computational sciences who have called for standards to include data and code with research findings, oriented towards the dissemination of reproducible research, facilitating innovation, and enabling broader communication of scientific outputs [137].

Data Contestability. Designing mechanisms of contestation into algorithmic decision making systems has been motivated as a way of encouraging critical and responsible engagement between various stakeholders, surfacing values that might not otherwise be visible, and safeguarding a system against misuse [68, 103]. However, the ability of a given stakeholder to challenge an automated decision is limited, in large part, by the degree to which different components of the algorithmic system are legible to those outside the development team. By surfacing otherwise invisible aspects of dataset development the proposed data documentation frameworks creates the conditions necessary for contesting datasets as well as algorithmic systems built from them.

From Metadata to Data. Information about processes of dataset development constitutes metadata, *i.e.*, structured data about data [48]. This metadata can in some cases be incorporated *into the dataset itself*, analogous to how digital image files record time, date and camera settings [10]. In the case of data labeling, this collapses the dichotomy of the the observer and the observed, making label provenance explicit in the data [136] rather than pretending to be a “view from nowhere” [65, 104] and creating the space for alternate epistemologies needed to confront the bias paradox [66]. Reconceiving dataset records as essentially relational artefacts, between the beholder and beheld, enables datasets to document aspects of their own creation. Making epistemological uncertainties in data explicit in the data itself somewhat parallels engineering cultures of exception-handling, in which classes of errors are thrown and caught (see Section 3), *i.e.*, it is the *unknown* errors which are the most dangerous in safety-critical applications.⁵

5.2 Ecologies of Data Work

Based on our analysis above of the processes needed to improve dataset development documentation, there are a number of ecological challenges that potentially inhibit process adoption. To mitigate these challenges, we propose the following goals for research, academia and industry.

Recognize AI dataset expertise. AI datasets take more effort and rigor to curate and manage than code does [4], and yet this work remains underappreciated. Omitting dataset development from the core training of ML practitioners is comparable to not teaching computer programmers about data structures. Following Jo and Gebru, we call for cultivating *AI Dataset Development* as a distinct subfield of AI (with parallels to archival studies and corpus linguistics) [75], providing scaffolding that serves as a primary and

⁵The Apollo missions of the 1960s “flew with program errors that were thoroughly documented, and the effects were understood before the mission. No *unknown* programming errors were discovered during a mission.” [64] (original emphasis).

fundamental input to the development lifecycle. Seen through this lens, dataset development requires a great deal of further dedicated research (e.g., on practices of collection, methods and protocols), its own theory [35],⁶ conferences, prizes, *etc.* There is a need to shift the perception of dataset work away from crafty gluework to highly skilled technical infrastructure work. Data provides a fundamental scaffolding from which further ML development can launch. Recasting AI dataset work from drudgework or necessary evil to valued careful cultivation and stewardship means acknowledging the roots of dataset work in feminized practice, notably informed by proximate fields in library and information studies [69]. Some seeds worth cultivating include values-driven active learning [5], diversity-informed data collection [136], and challenge data sets [39].⁷

Engage with data disciplines beyond ML. Building on our call for engagement with subject domain experts when building datasets (Section 4.2.2), we also call for engagement with disciplines with rich dataset development practices. We echo calls to learn from archival studies [75] and language corpus design [150], and to pay greater attention to research in human computation [141]. AI Dataset Development could also benefit from engagement with established data quality and data integrity practices from clinical research. In medicine data integrity refers to the completeness, consistency, and accuracy of data, and to achieve these data should be attributable, legible, contemporaneously recorded, original or a true copy, and accurate (ALCOA) [47]. The US FDA states that “System design and controls should enable easy detection of errors, omissions, and aberrant results throughout the data’s life cycle” [48]. What could analogous practices look like for AI dataset development?

Understand cultures of dataset development. There is need for more research on the cultures of AI Dataset Development, echoing work on cultures of computer programming (see Section 3) and cultures of mathematics [95]. Recent work in this direction includes Denton and Hanna *et al.*, who discuss focusing on “bringing the people back in” to the discussion of dataset construction [41]. Better understanding is needed of how the deeply collaborative nature of much dataset work conflicts with “disincentives to collect data that come from a system that emphasizes individual academic output” [87], as well as relationships with other common ML myopias—especially common to the “hacker culture” of ML (Section 3)—that celebrate Alpha leadership, competition, and other masculine tropes.⁸

Invest in organizational processes and structures. Adopt best practices from managerial theory to software development, including critically the needs for properly budgeting for maintenance and treating it properly as “repair work” [124], for redressing dataset technical debt, and for organizational recognition and rewards for data work. We need to shift processes such that explicit normative considerations are no longer rarely in mind [109].

⁶Especially since academia values theory over practice [57, 61].

⁷These currently seem most common in NLP, e.g., [78, 93, 102, 118, 148, 151, 152, 154].

⁸The connection of “engineering” itself with computer programming was an intentionally, highly gendered endeavor intended to shift work from a primarily feminized craft to something that would pay more and attract men [1].

Demand greater methodological clarity of ML research. Whereas methods are tools of research, *methodologies* are commonly described as the principles, values or frames of reference of research [98]. The methodologies of ML research are often opaque or obscure, e.g., is the goal to acquire *knowledge-how* or *knowledge-that* (Section 3)? Is the research akin to mathematics, science, applied science, engineering, engineering science, or something else? These questions matter not just because they bring justification to methods (why are robustness, AUC, effect sizes or *p*-values the relevant things to measure?), but also because the ML goals determine the ML dataset requirements. ML datasets need to be understood as engineered artefacts in support of ML goals; if we instead mistake dataset development for “academic science”, we perpetuate dangerous assumptions about objectivity.

6 CONCLUSIONS

We have described the role of data in AI system development and highlighted the gap between current development processes and processes that can address modern AI outcomes. We ground development within concrete ethical practices, proposing frameworks to operationalize transparency and accountability. Successful implementation of our proposed frameworks depends critically on clearly formulated development processes, with discrete and interconnected stages. Owners charged with responsibility for each stage must work by consulting appropriate experts and produce detailed accounts of what happens during each stage, and why. The form of these accounts (requirements specifications, design documents, testing reports, *etc.*) actively work to improve the quality, reliability and validity of the data itself.

Our approach is informed by the software development lifecycle, based on two correspondences that we detailed above: 1) between data and technical infrastructure, and 2) between the processes and goals of engineering with those of dataset development. Based on these similarities, we presented a clear and comprehensive framework that builds on existing best practices. We leverage methods in related fields to describe a procedure for documenting all stages of dataset development: requirements analysis, design, implementation, evaluation, and maintenance. The detailed bookkeeping of accounts of each stage of the lifecycle requires owners with distinct knowledge, skills, and processes, and the expertise most critical to consult at each stage is similarly distinct: stakeholders and impacted parties; subject matter experts; HCI, data operations managers and raters with domain expertise; data scientists; supporters of tooling and ongoing funding.

We argue that frameworks for transparency and accountability are badly needed, as systemic devaluation of dataset work has contributed greatly to its historic failures. We believe that a culture shift is needed to truly embrace the modernization of AI development processes towards ethically-informed frameworks. We advocate for valuing work on data commensurate with the value placed on technical infrastructure work and model building. If we are to improve accountability for the datasets that power ML, then, perhaps paradoxically, we need to create avenues for which datasets can earn accolades when they are properly due, avenues which encourage skilful data experts to proudly say “I am responsible for this.”

REFERENCES

- [1] Janet Abbate. 2012. *Recoding gender: Women's changing participation in computing*. MIT Press.
- [2] Alain Abran, James W Moore, Pierre Bourque, Robert Dupuis, and L Tripp. 2004. Software engineering body of knowledge. *IEEE Computer Society, Angela Burgess* (2004).
- [3] Irfan A Alvi. 2013. Engineers need to get real, but can't: The role of models. In *Structures Congress 2013: Bridging Your Passion with Your Profession*. 916–927.
- [4] Saleema Amershi, Andrew Begel, Christian Bird, Robert DeLine, Harald Gall, Ece Kamar, Nachiappan Nagappan, Besmira Nushi, and Thomas Zimmermann. 2019. Software engineering for machine learning: A case study. In *2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*. IEEE, 291–300.
- [5] Hadis Anahideh and Abolfazl Asudeh. 2020. Fair Active Learning. *arXiv preprint arXiv:2001.01796* (2020).
- [6] Helmut K Anheier. 2017. Infrastructure and the principle of the hiding hand. *The Governance of Infrastructure* (2017), 63.
- [7] Itamar Arel. 2012. Deep reinforcement learning as foundation for artificial general intelligence. In *Theoretical Foundations of Artificial General Intelligence*. Springer, 89–102.
- [8] Matthew Arnold, Rachel KE Bellamy, Michael Hind, Stephanie Houde, Sameep Mehta, A Mojsilović, Ravi Nair, K Natesan Ramamurthy, Alexandra Olteanu, David Piorkowski, et al. 2019. FactSheets: Increasing trust in AI services through supplier's declarations of conformity. *IBM Journal of Research and Development* 63, 4/5 (2019), 6–1.
- [9] Rob Ashmore, Radu Calinescu, and Colin Paterson. 2019. Assuring the machine learning lifecycle: Desiderata, methods, and challenges. *arXiv preprint arXiv:1905.04223*.
- [10] Camera & Imaging Products Association et al. 2010. Exchangeable image file format for digital still cameras: Exif Version 2.3. *CIPA DC-008 Translation-2010* (2010).
- [11] Marzieh Babaeianjelodar, Stephen Lorenz, Josh Gordon, Jeanna Matthews, and Evan Freitag. 2020. Quantifying Gender Bias in Different Corpora. In *Companion Proceedings of the Web Conference 2020*. 752–759.
- [12] Iain Barclay, Alun Preece, Ian Taylor, and Dinesh Verma. 2019. Towards Traceability in Data Ecosystems using a Bill of Materials Model. *arXiv preprint arXiv:1904.04253* (2019).
- [13] Solon Barocas and Andrew D Selbst. 2016. Big data's disparate impact. *Calif. L. Rev.* 104 (2016), 671.
- [14] Denis Baylor, Eric Breck, Heng-Tze Cheng, Noah Fiedel, Chuan Yu Foo, Zakaria Haque, Salem Haykal, Mustafa Ispir, Vihan Jain, Levent Koc, et al. 2017. TFX: A tensorflow-based production-scale machine learning platform. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1387–1395.
- [15] Peter M Bednar and Christine Welch. 2009. Contextual inquiry and requirements shaping. In *Information Systems Development*. Springer, 225–236.
- [16] Emily M Bender and Batya Friedman. 2018. Data statements for natural language processing: Toward mitigating system bias and enabling better science. *Transactions of the Association for Computational Linguistics* 6 (2018), 587–604.
- [17] Keith H Bennett and Václav T Rajlich. 2000. Software maintenance and evolution: a roadmap. In *Proceedings of the Conference on the Future of Software Engineering*. 73–87.
- [18] Elena Beretta, Antonio Vetrò, Bruno Lepri, and Juan Carlos De Martin. 2018. Ethical and Socially-Aware Data Labels. In *Annual International Symposium on Information Management and Big Data*. Springer, 320–327.
- [19] Anandhi S Bharadwaj. 2000. A resource-based perspective on information technology capability and firm performance: an empirical investigation. *MIS quarterly* (2000), 169–196.
- [20] Douglas Biber. 1993. Representativeness in corpus design. *Literary and linguistic computing* 8, 4 (1993), 243–257.
- [21] Matthias Boehm, Iulian Antonov, Sebastian Baunsgaard, Mark Dokter, Robert Ginthör, Kevin Innerebner, Florijan Klezin, Stefanie Lindstaedt, Arnab Phani, Benjamin Rath, Berthold Reinwald, Shafiq Siddiqi, and Sebastian Benjamin Wrede. 2019. SystemDS: A Declarative Machine Learning System for the End-to-End Data Science Lifecycle. *arXiv preprint arXiv:1909.02976*.
- [22] Geoffrey C Bowker and Susan Leigh Star. 2000. *Sorting things out: Classification and its consequences*. MIT press.
- [23] Eric Breck, Shanjing Cai, Eric Nielsen, Michael Salib, and D Sculley. 2017. The ML test score: A rubric for ML production readiness and technical debt reduction. In *2017 IEEE International Conference on Big Data (Big Data)*. IEEE, 1123–1132.
- [24] Harvey Brooks. 1994. The relationship between science and technology. *Research policy* 23, 5 (1994), 477–486.
- [25] Richard Buchanan. 1992. Wicked problems in design thinking. *Design issues* 8, 2 (1992), 5–21.
- [26] William Bulleit, Jon Schmidt, Irfan Alvi, Erik Nelson, and Tonatiuh Rodriguez-Nikl. 2015. Philosophy of engineering: What it is and why it matters. *Journal of Professional Issues in Engineering Education and Practice* 141, 3 (2015), 02514003.
- [27] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. 77–91.
- [28] Kaylee Burns, Lisa Hendricks, Trevor Darrell, and Anna Rohrbach. 2018. Women also Snowboard: Overcoming Bias in Captioning Models. (03 2018).
- [29] Lawrence Busch. 2014. A dozen ways to get lost in translation: Inherent challenges in large scale data sets. *International Journal of Communication* 8 (2014), 18.
- [30] Luiz Fernando Capretz, Daniel Varona, and Arif Raza. 2015. Influence of personality types in software tasks choices. *Computers in Human behavior* 52 (2015), 373–378.
- [31] Supriyo Chakraborty, Richard Tomsett, Ramya Raghavendra, Daniel Harborne, Moustafa Alzantot, Federico Cerutti, Mani Srivastava, Alun Preece, Simon Julier, Raghuvier M Rao, et al. 2017. Interpretability of deep learning models: a survey of results. In *2017 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCOM/ITP/SCI)*. IEEE, 1–6.
- [32] Danielle Citron and Frank Pasquale. 2014. The scored society: Due process for automated predictions. *Washington Law Review* 89 (03 2014), 1–33.
- [33] CMMI Institute. 2018. Patient Demographic Data Quality (PDDQ) Framework.
- [34] IEEE Standards Coordinating Committee et al. 1990. IEEE Standard Glossary of Software Engineering Terminology (IEEE Std 610.12-1990). Los Alamitos, CA: IEEE Computer Society 169 (1990).
- [35] Peter V Coveney, Edward R Dougherty, and Roger R Highfield. 2016. Big data need big theory too. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 374, 2080 (2016), 20160153.
- [36] Kate Crawford, Mary L Gray, and Kate Miltner. 2014. Critiquing Big Data: Politics, ethics, epistemology (special section introduction). *International Journal of Communication* 8 (2014), 10.
- [37] Kate Crawford and Trevor Paglen. 2019. Excavating AI: The politics of images in machine learning training sets. *Excavating AI* (2019).
- [38] Thomas Davidson, Debasmita Bhattacharya, and Ingmar Weber. 2019. Racial bias in hate speech and abusive language detection datasets. *arXiv preprint arXiv:1905.12516* (2019).
- [39] Ernest Davis. 2014. The limitations of standardized science tests as benchmarks for artificial intelligence research: Position paper. *arXiv preprint arXiv:1411.1629* (2014).
- [40] Terrance de Vries, Ishan Misra, Changhan Wang, and Laurens van der Maaten. 2019. Does object recognition work for everyone?. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 52–59.
- [41] Emily Denton, Alex Hanna, Razvan Amironesei, Andrew Smart, Hilary Nicole, and Morgan Klaus Scheuerman. 2020. Bringing the People Back In: Contesting Benchmark Machine Learning Datasets. *arXiv preprint arXiv:2007.07399* (2020).
- [42] Torgeir Dingsøyr. 2005. Postmortem reviews: purpose and approaches in software engineering. *Information and Software Technology* 47, 5 (2005), 293–303.
- [43] Lisa Ehrlinger, Elisa Rusz, and Wolfram Wöfl. 2019. A survey of data quality measurement and monitoring tools. *arXiv preprint arXiv:1907.08138* (2019).
- [44] Virginia Eubanks. 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.
- [45] Heinz Eulau. 2007. *Crossroads of social science: The ICPSR 25th anniversary volume*. Algora Publishing.
- [46] Bent Flyvbjerg and Cass R Sunstein. 2016. The principle of the malevolent hiding hand; or, the planning fallacy writ large. *Social Research: An International Quarterly* 83, 4 (2016), 979–1004.
- [47] Food, Drug Administration, et al. 2016. Data Integrity and Compliance With CGMP Guidance for Industry. *Draft Guidance* (2016).
- [48] US Food, Drug Administration, et al. 2018. Data Integrity and Compliance with Drug CGMP Questions and Answers Guidance for Industry.
- [49] Brett M Frischmann. 2012. *Infrastructure: The social value of shared resources*. Oxford University Press.
- [50] Jonathan Furner. 2016. "Data": The data. In *Information cultures in the digital age*. Springer, 287–306.
- [51] Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences* 115, 16 (2018), E3635–E3644.
- [52] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2018. Datasheets for datasets. *arXiv preprint arXiv:1803.09010* (2018).
- [53] R Stuart Geiger, Kevin Yu, Yanlai Yang, Mindy Dai, Jie Qiu, Rebekah Tang, and Jenny Huang. 2020. Garbage in, garbage out? do machine learning application papers in social computing report where human-labeled training data comes from?. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 325–336.
- [54] Dave Gershgorin. 2018. If AI is going to be the world's doctor, it needs better textbooks. <https://qz.com/1367177/if-ai-is-going-to-be-the-worlds-doctor-it-needs-better-textbooks>. Accessed: 2020-09-19.
- [55] Rick O Gilmore, Karen E Adolph, and David S Millman. 2016. Curating identifiable data for sharing: The databrary project. In *2016 New York Scientific Data*

- Summit (NYSYS). IEEE, 1–6. <https://doi.org/10.1109/NYSYS.2016.7747817>
- [56] Martin Glinz. 2011. A glossary of requirements engineering terminology. *Standard Glossary of the Certified Professional for Requirements Engineering (CPRE) Studies and Exam, Version 1* (2011), 56.
- [57] Steven L Goldman. 2004. Why we need a philosophy of engineering: a work in progress. *Interdisciplinary Science Reviews* 29, 2 (2004), 163–176.
- [58] Steven L Goldman. 2010. Beyond satisficing: Design, trade offs and the rationality of engineering. In *2010 Forum on philosophy, engineering & technology*.
- [59] Ben Green. 2020. Data science as political action: Grounding data science in a politics of justice. Available at SSRN 3658431 (2020).
- [60] Tristan Greene. 2020. 2010–2019: The rise of deep learning. <https://thenextweb.com/artificial-intelligence/2020/01/02/2010-2019-the-rise-of-deep-learning/>. Accessed: 2020-09-26.
- [61] Ian Hacking, Jan Hacking, et al. 1983. *Representing and intervening: Introductory topics in the philosophy of natural science*. Cambridge University Press.
- [62] Alon Halevy, Peter Norvig, and Fernando Pereira. 2009. The unreasonable effectiveness of data. *IEEE Intelligent Systems* 24, 2 (2009), 8–12.
- [63] Brendan Hall and Kevin Driscoll. 2014. Distributed System Design Checklist. (2014).
- [64] Eldon C Hall. 1996. *Journey to the moon: the history of the Apollo guidance computer*. Aiaa.
- [65] Donna Haraway. 1988. Situated knowledges: The science question in feminism and the privilege of partial perspective. *Feminist studies* 14, 3 (1988), 575–599.
- [66] Deborah K Heikes. 2004. The bias paradox: why it's not just for feminists anymore. *Synthese* 138, 3 (2004), 315–335.
- [67] Benjamin Heinzerling. 2020. NLP's Clever Hans Moment has Arrived. *Journal of Cognitive Science* 21, 1 (2020), 159–167.
- [68] Tad Hirsch, Kritzia Merced, Shrikanth Narayanan, Zac E Imel, and David C Atkins. 2017. Designing contestability: Interaction design, machine learning, and mental health. In *Proceedings of the 2017 Conference on Designing Interactive Systems*. 95–99.
- [69] Anna Lauren Hoffmann and Raina Bloom. 2016. Digitizing Books, Obscuring Women's Work: Google Books, Librarians, and Ideologies of Access. *Ada: A Journal of Gender, New Media, and Technology* 9 (2016).
- [70] Sarah Holland, Ahmed Hosny, Sarah Newman, Joshua Joseph, and Kasia Chmielinski. 2018. The dataset nutrition label: A framework to drive higher data quality standards. *arXiv preprint arXiv:1805.03677* (2018).
- [71] Kenneth Holstein, Jennifer Wortman Vaughan, Hal Daumé III, Miro Dudik, and Hanna Wallach. 2019. Improving fairness in machine learning systems: What do industry practitioners need?. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [72] Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton, Kellie Webster, Yu Zhong, and Stephen Denuyl. 2020. Social Biases in NLP Models as Barriers for Persons with Disabilities. *ACL* (2020).
- [73] Lilly C Irani and M Six Silberman. 2013. Turkopticon: Interrupting worker invisibility in amazon mechanical turk. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 611–620.
- [74] Abigail Z Jacobs and Hanna Wallach. 2019. Measurement and fairness. *arXiv preprint arXiv:1912.05511* (2019).
- [75] Eun Seo Jo and Timnit Gebru. 2020. Lessons from archives: strategies for collecting sociocultural data in machine learning. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 306–316.
- [76] Jacint Jordana. 2017. Accountability Challenges in the Governance of Infrastructure. *The Governance of Infrastructure* (2017), 43.
- [77] Taehee Jung, Dongyeop Kang, Lucas Mentch, and Eduard Hovy. 2019. Earlier Isn't Always Better: Sub-aspect Analysis on Corpus and System Biases in Summarization. *arXiv preprint arXiv:1908.11723* (2019).
- [78] Daniel Khashabi, Snigdha Chaturvedi, Michael Roth, Shyam Upadhyay, and Dan Roth. 2018. Looking beyond the surface: A challenge set for reading comprehension over multiple sentences. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. 252–262.
- [79] Adam Kilgariff and Gregory Grefenstette. 2003. Introduction to the special issue on the web as corpus. *Computational linguistics* 29, 3 (2003), 333–347.
- [80] Gary King. 2007. An introduction to the dataverse network as an infrastructure for data sharing.
- [81] Gary Klein. 2007. Performing a project premortem. *Harvard business review* 85, 9 (2007), 18–19.
- [82] Billy Vaughn Koen. 2003. Discussion of The Method: Conducting the Engineer's Approach to Problem Solving.
- [83] Nitin Kohli, Renata Barreto, and Joshua A Kroll. 2018. Translation tutorial: a shared lexicon for research and practice in human-centered software systems. In *1st Conference on Fairness, Accountability, and Transparency*. New York, NY, USA, Vol. 7.
- [84] Tobias Krafft, Marc Hauer, Lajla Fetic, Andreas Kaminski, Michael Puntschuh, Philipp Otto, Christoph Hubig, Torsten Fleischer, Paul Grünke, Rafaela Hillerbrand, Carla Hustedt, and Sebastian Hallensleben. 2020. From Principles to Practice - An interdisciplinary framework to operationalise AI ethics.
- [85] Jonathan Krause, Hailin Jin, Jianchao Yang, and Li Fei-Fei. 2015. Fine-grained recognition without part annotations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5546–5555.
- [86] Hiroshi Kuwajima, Hirotohi Yasuoka, and Toshihiro Nakae. 2020. Engineering problems in machine learning systems. *Machine Learning* (2020), 1–24.
- [87] Stuart N Lane. 2020. Editorial 2020 Part II: Data from nowhere? *Earth Surface Processes and Landforms* 45, 1 (2020), 5–10.
- [88] Brian Larkin. 2013. The politics and poetics of infrastructure. *Annual review of anthropology* 42 (2013), 327–343.
- [89] Bruno Latour. 1987. *Science in action: How to follow scientists and engineers through society*. Harvard university press.
- [90] Bruno Latour and Steve Woolgar. 2013. *Laboratory life: The construction of scientific facts*. Princeton University Press.
- [91] Alexander Lavin. 2020. Machine Learning Is No Place To "Move Fast And Break Things". <https://www.forbes.com/sites/alexanderlavin/2020/02/17/machine-learning-is-no-place-to-move-fast-and-break-things/#2bfe96d15f2>. Accessed: 2020-09-19.
- [92] Sabina Leonelli. 2020. Scientific Research and Big Data. <https://plato.stanford.edu/entries/science-big-data/>. Accessed: 2020-10-04.
- [93] Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. 2020. LogQA: A Challenge Dataset for Machine Reading Comprehension with Logical Reasoning. *arXiv preprint arXiv:2007.08124* (2020).
- [94] Xiaoxuan Liu, Livia Faes, Aditya U Kale, Siegfried K Wagner, Dun Jack Fu, Alice Bruynseels, Thushika Mahendiran, Gabriella Moraes, Mohith Shandas, Christoph Kern, et al. 2019. A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. *The lancet digital health* 1, 6 (2019), e271–e297.
- [95] Eric Livingston. 1999. Cultures of proving. *Social studies of science* 29, 6 (1999), 867–888.
- [96] David Loshin. 2010. *Master data management*. Morgan Kaufmann.
- [97] David Loshin. 2010. *The practitioner's guide to data quality improvement*. Elsevier.
- [98] Noella Mackenzie and Sally Knipe. 2006. Research dilemmas: Paradigms, methods and methodology. *Issues in educational research* 16, 2 (2006), 193–205.
- [99] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*. 220–229.
- [100] Brent Daniel Mittelstadt, Patrick Allo, Mariarosaria Taddeo, Sandra Wachter, and Luciano Floridi. 2016. The ethics of algorithms: Mapping the debate. *Big Data & Society* 3, 2 (2016), 2053951716679679.
- [101] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing Atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [102] James Mullenbach, Jonathan Gordon, Nanyun Peng, and Jonathan May. 2019. Do Nuclear Submarines Have Nuclear Captains? A Challenge Dataset for Commonsense Reasoning over Adjectives and Objects. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 6054–6060.
- [103] Deirdre K Mulligan, Daniel Kluttz, and Nitin Kohli. 2019. Shaping Our Tools: Contestability as a Means to Promote Responsible Algorithmic Decision Making in the Professions. Available at SSRN 3311894 (2019).
- [104] Thomas Nagel. 1989. *The view from nowhere*. Oxford University Press.
- [105] Gina Neff, Anissa Tanweer, Brittany Fiore-Gartland, and Laura Osburn. 2017. Critique and contribute: A practice-based framework for improving critical data studies and data science. *Big data* 5, 2 (2017), 85–97.
- [106] Helen Nissenbaum. 1996. Accountability in a computerized society. *Science and engineering ethics* 2, 1 (1996), 25–42.
- [107] Roy Oshrover. 2015. *The art of unit testing*. MITP-Verlags GmbH & Co. KG.
- [108] Irene V Pasquetto, Bernadette M Randles, and Christine L Borgman. 2017. On the reuse of scientific data. *Data Science Journal* 16 (2017), 8.
- [109] Samir Passi and Solon Barocas. 2019. Problem formulation and fairness. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 39–48.
- [110] Ron Patton. 2006. *Software testing*. Pearson Education India.
- [111] Tomas Petricek. 2019. Cultures of programming. (2019). unpublished.
- [112] Heather A Piwowar and Todd J Vision. 2013. Data reuse and the open data citation advantage. *PeerJ* 1 (2013), e175.
- [113] Neoklis Polyzotis, Sudip Roy, Steven Euijong Whang, and Martin Zinkevich. 2018. Data lifecycle challenges in production machine learning: a survey. *ACM SIGMOD Record* 47, no. 2 (2018), 17–28.
- [114] Vinay Uday Prabhu and Abeba Birhane. 2020. Large image datasets: A pyrrhic win for computer vision? *arXiv preprint arXiv:2006.16923* (2020).
- [115] Inioluwa Deborah Raji, Andrew Smart, Rebecca N White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI accountability gap: defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on*

- Fairness, Accountability, and Transparency*. 33–44.
- [116] Inioluwa Deborah Raji and Jingying Yang. 2019. ABOUT ML: Annotation and Benchmarking on Understanding and Transparency of Machine Learning Lifecycles. *arXiv preprint arXiv:1912.06166* (2019).
 - [117] Ari Ramkilwan. 2018. The rise and rise of AI in Africa. <https://medium.com/datadriveninvestor/the-rise-and-rise-of-ai-in-africa-a6cf6bf89217>. Accessed: 2020-09-26.
 - [118] Matthew Richardson, Christopher J.C. Burges, and Erin Renshaw. 2013. MCTest: A Challenge Dataset for the Open-Domain Machine Comprehension of Text. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Seattle, Washington, USA, 193–203. <https://www.aclweb.org/anthology/D13-1020>
 - [119] Rashida Richardson, Jason M Schultz, and Kate Crawford. 2019. Dirty data, bad predictions: How civil rights violations impact police data, predictive policing systems, and justice. *NYUJ Rev. Online* 94 (2019), 15.
 - [120] Horst WJ Rittel and Melvin M Webber. 1973. Dilemmas in a general theory of planning. *Policy sciences* 4, 2 (1973), 155–169.
 - [121] Yuji Roh, Geon Heo, and Steven Euijong Whang. 2019. A survey on data collection for machine learning: a big data-AI integration perspective. *IEEE Transactions on Knowledge and Data Engineering* (2019).
 - [122] Jeanne W Ross, Cynthia Mathis Beath, and Dale L Goodhue. 1996. Develop long-term competitiveness through IT assets. *Sloan management review* 38, 1 (1996), 31–42.
 - [123] Gilbert Ryle. 1945. Knowing how and knowing that: The presidential address. In *Proceedings of the Aristotelian society*, Vol. 46. JSTOR, 1–16.
 - [124] SE Sachs. 2019. The algorithm at work? Explanation and repair in the enactment of similarity in art data. *Information, Communication & Society* (2019), 1–17.
 - [125] Nithya Sambasivan, Diana Akrong, Hannah Highfill, Lora Mois Aroyo, Praveen Kumar Paritosh, and Shivani Kapania. 2021. "Everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI. In *Proceedings of CHI 2021*.
 - [126] Morgan Klaus Scheuerman, Kandrea Wade, Caitlin Lustig, and Jed R. Brubaker. 2020. How We've Taught Algorithms to See Identity: Constructing Race and Gender in Image Databases for Facial Analysis. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1, Article 058 (May 2020), 35 pages. <https://doi.org/10.1145/3392866>
 - [127] David Sculley, Gary Holt, Daniel Golovin, Eugene Davydov, Todd Phillips, Dietmar Ebner, Vinay Chaudhary, Michael Young, Jean-Francois Crespo, and Dan Dennison. 2015. Hidden technical debt in machine learning systems. In *Advances in neural information processing systems*. 2503–2511.
 - [128] David Sculley, Jasper Snoek, Alex Wiltschko, and Ali Rahimi. 2018. Winner's curse? On pace, progress, and empirical rigor. (2018).
 - [129] Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 59–68.
 - [130] Burr Settles. 2009. *Active learning literature survey*. Technical Report. University of Wisconsin-Madison Department of Computer Sciences.
 - [131] Shreya Shankar, Yoni Halpern, Eric Breck, James Atwood, Jimbo Wilson, and D Sculley. 2017. No classification without representation: Assessing geodiversity issues in open data sets for the developing world. *arXiv preprint arXiv:1711.08536* (2017).
 - [132] David Shoemaker. 2011. Attributability, answerability, and accountability: Toward a wider theory of moral responsibility. *Ethics* 121, 3 (2011), 602–632.
 - [133] David Shoemaker. 2015. *Responsibility from the Margins*. Oxford University Press, USA.
 - [134] Herbert A Simon. 1973. The structure of ill structured problems. *Artificial intelligence* 4, 3-4 (1973), 181–201.
 - [135] Andrew Smart, Larry James, Ben Hutchinson, Simone Wu, and Shannon Vallor. 2020. Why Reliabilism Is not Enough: Epistemic and Moral Justification in Machine Learning. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 372–377.
 - [136] Katherine Stasaski, Grace Hui Yang, and Marti A Hearst. 2020. More Diverse Dialogue Datasets via Diversity-Informed Data Collection. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 4958–4968.
 - [137] Victoria Stodden and Sheila Miguez. 2014. Best Practices for Computational Science: Software Infrastructure and Environments for Reproducible and Extensible Research. *Journal of Open Research Software* 2, 1 (2014).
 - [138] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. 2017. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision*. 843–852.
 - [139] Herb Sutter and Andrei Alexandrescu. 2004. *C++ coding standards: 101 rules, guidelines, and best practices*. Pearson Education.
 - [140] Astra Taylor. 2018. The automation charade. *Logic Magazine* (2018).
 - [141] Jennifer Wortman Vaughan. 2018. Making Better Use of the Crowd: How Crowdsourcing Can Advance Machine Learning Research. *Journal of Machine Learning Research* 18, 193 (2018), 1–46. <http://jmlr.org/papers/v18/17-234.html>
 - [142] Andreas Vogelsang and Markus Borg. 2019. Requirements Engineering for Machine Learning: Perspectives from Data Scientists. In *2019 IEEE 27th International Requirements Engineering Conference Workshops (REW)*. IEEE, 245–251.
 - [143] Joel Walmsley. 2020. Artificial intelligence and the value of transparency. *AI & SOCIETY* (2020), 1–11.
 - [144] Kai Wegrich and Gerhard Hammerschmid. 2017. Infrastructure governance as political choice. *The governance of infrastructure* (2017), 21–42.
 - [145] Michael Weisberg. 2012. *Simulation and similarity: Using models to understand the world*. Oxford University Press.
 - [146] Chris Welty, Praveen Paritosh, and Lora Aroyo. 2019. Metrology for AI: From Benchmarks to Instruments. *arXiv preprint arXiv:1911.01875* (2019).
 - [147] Maranke Wieringa. 2020. What to account for when accounting for algorithms: a systematic literature review on algorithmic accountability. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 1–18.
 - [148] Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. Association for Computational Linguistics, New Orleans, Louisiana, 1112–1122. <https://doi.org/10.18653/v1/N18-1101>
 - [149] Langdon Winner. 1980. Do artifacts have politics? *Daedalus* (1980), 121–136.
 - [150] Richard Xiao. 2010. Corpus creation. *Handbook of Natural Language Processing (2n Revised edition)* (2010), 147–165.
 - [151] Semih Yagcioglu, Aykut Erdem, Erkan Erdem, and Nazli Ilkizler-Cimbis. 2018. RecipeQA: A challenge dataset for multimodal comprehension of cooking recipes. *arXiv preprint arXiv:1809.00812* (2018).
 - [152] Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. WikiQA: A Challenge Dataset for Open-Domain Question Answering. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Lisbon, Portugal, 2013–2018. <https://doi.org/10.18653/v1/D15-1237>
 - [153] Matei Zaharia, Andrew Chen, Aaron Davidson, Ali Ghodsi, Sue Ann Hong, Andy Konwinski, Siddharth Murching, Tomas Nykodym, Paul Ogilvie, Mani Parkhe, Fen Xie, and Corey Zumaret. 2018. Accelerating the Machine Learning Lifecycle with MLflow. In *IEEE Data Eng. Bull.* 41, no. 4. 39–45.
 - [154] Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. Swag: A large-scale adversarial dataset for grounded commonsense inference. *arXiv preprint arXiv:1808.05326* (2018).
 - [155] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2017. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. *arXiv preprint arXiv:1707.09457* (2017).
 - [156] Junji Zhi, Vahid Garousi-Yusifoglu, Bo Sun, Golar Garousi, Shawn Shahnewaz, and Guenther Ruhe. 2015. Cost, benefits and quality of software development documentation: A systematic mapping. *Journal of Systems and Software* 99 (2015), 175–198.
 - [157] Matthew Zook, Solon Barocas, Danah Boyd, Kate Crawford, Emily Keller, Seeta Peña Gangadharan, Alyssa Goodman, Rachelle Hollander, Barbara A Koenig, Jacob Metcalf, et al. 2017. Ten simple rules for responsible big data research.

APPENDIX A: TEMPLATE FOR DATASET REQUIREMENTS SPECIFICATION

Name of Dataset: Requirements Specification

Owner: *Name*; Created: *Date*; Last updated: *Date*

Vision

Brief summary of the envisioned data(set), its domains and scope.

Motivation

Problem and context that motivate why the data is needed.

Intended uses

Specific uses of the data that are intended.

Non-intended uses

What is the data not intended for? What should the data not be used for, and why?

Glossary of terms

If relevant, brief summary of acronyms and domain specific concepts for the general reader.

Related documents

List any related documents.

Data mocks

Include 2-3 typical examples of what the data instances should "look" like.

Stakeholders consulted

Whose needs were consulted and synthesised when creating this document? How were conflicting needs resolved?

Creation requirements

Where should the data come from? Include sources and collection methods

- *Name of the requirement. Description.*
- *Name of the requirement. Description.*

Instance requirements

What requirements are there for data instances? Include any acceptable tradeoffs. Include numbers and types of instances, features, and labels.

- *Name of the requirement. Description.*
- *Name of the requirement. Description.*

Distributional requirements

What requirements are there for the distributions of your data? Include any acceptable tradeoffs. Include sampling requirements. If your data represents a set of people, describe who should be represented and in what numbers.

- *Name of the requirement. Description.*
- *Name of the requirement. Description.*

Data processing requirements

How should the data be annotated and filtered? Who should do the annotating? How should data be validated? Include any acceptable tradeoffs.

- *Name of the requirement. Description.*
- *Name of the requirement. Description.*

Performance requirements

What can people who use this dataset for its intended uses expect?

- *Name of the requirement. Description.*
- *Name of the requirement. Description.*

Maintenance requirements

Should the data be regularly updated? If so, how often? For how long should the data be retained? Include any acceptable tradeoffs.

Sharing requirements

Should the data be made available to other teams within Google and/or open-sourced? If so, what constraints on data licensing, access, usage, and distribution are needed? Include any acceptable tradeoffs.

Caveats and risks

What would be the consequences of using data meeting the requirements described above?

Data ethics

Document your considerations of the ethical implications of the data and its collection.

Sign-off grid

Name	Role	Date

Changelog

Editor	Comments	Date

APPENDIX B: TEMPLATE FOR DATASET DESIGN DOCUMENT

When adapting the template, we recommend reflecting on questions such as “If you were asked to review the design of this dataset, what questions would you ask?” and “If you were presenting a design, what questions would you dread being asked?” [63].

Name of Dataset: Design Document

Owner: *Name*; Created: *Date*; Last updated: *Date*

Overview

High-level overview of the dataset.

Dataset Name: Name of dataset.

Primary Data Type(s): Primary data types; Eg: images, video, text.

Data Content: Eg. bounding boxes, image labels

Funding: How was the dataset funded?

Objective

What are the key objectives of the dataset? Is there a requirements specification?

Version

Current version; Differences to previous versions.

Background

Describe any relevant background of the dataset

Sources

System details; Where will the data come from? Selection and sampling criteria.

Annotations

Features and labels; Who are the annotators? How will they be trained?

Ratings: Rating tasks; Rating types; Rating procedures;

Data Quality

How is quality measured? How are metrics validated?

Characteristics

Characteristics of the dataset.

Expected Characteristics: Eg. How many instances, features, ratings.

Correlations: Acceptable correlations; Unacceptable correlations;

Acceptable and Unacceptable Conjunctive Datasets: Datasets that can and can not be used in conjunction with this dataset?

Population: Population represented.

Privacy Handling

How is privacy handled?

Maintenance

Who will maintain the dataset? Is there a maintenance plan? What are the recovery strategies if issues arise?

Sharing

Will the dataset be shared? How will access be controlled? How will the dataset be licensed?

Caveats

Describe known caveats

Data Ethics

Ethical considerations; Mitigation.

Work estimates

How much time will it take to collect the data; Costs are involved.

Related Datasets

Which existing datasets are related to this one? Why are they unsuitable?

Dataset Discovery Process: How did you search for other datasets?

Survey: High level overview of related datasets.

	Your Dataset	Other Dataset 1	Other Dataset 2
Documentation and DOI	Yours	Theirs	Theirs
Motivation and Intended use	Yours	Theirs	Theirs
Location	Yours	Theirs	Theirs
Size, Sampling and Filtering	Yours	Theirs	Theirs
Annotation and Labels	Yours	Theirs	Theirs
(Expected) Performance	How will your dataset improve on existing ones?	How does it not meet your requirements?	How does it not meet your requirements?
Examples	Example instance	Example instance	Example instance

APPENDIX C: TEMPLATE FOR DATASET TESTING REPORT*Name of Dataset: Testing Report*Owner: *Name*; Created: *Date*; Last updated: *Date***Summary**

What is being tested?

Link to requirements specification.

Link to design document.

Meta-Testing

Is the data still needed? Are the data requirements still relevant and up-to-date?

Requirements Testing

Requirement tested	Results	Artifact
Requirement from requirements specification	Score or Results	Justification of the results or a link to artifact
Requirement from requirements specification	Score or Results	Justification of the results or a link to artifact
...	...	...

Untested Requirements

Untested Requirement	Reason for not testing
Requirement from requirements specification	Reason for not testing
Requirement from requirements specification	Reason for not testing
...	...

Adversarial Testing

Adversarial test	Results	Artifact
Describe tests	Score or Results	Justification of the results or a link to artifact
Describe tests	Score or Results	Justification of the results or a link to artifact
...	...	...

Other Testing

Test	Results	Artifact
Describe tests	Score or Results	Justification of the results or a link to artifact
Describe tests	Score or Results	Justification of the results or a link to artifact
...	...	...