

Accountability in an Algorithmic Society: Relationality, Responsibility, and Robustness in Machine Learning

A. FEDER COOPER, Cornell University, USA
 BENJAMIN LAUFER, Cornell Tech, USA
 EMANUEL MOSS, Cornell Tech and Data & Society Research Institute, USA
 HELEN NISSENBAUM, Cornell Tech, USA

In 1996, philosopher Helen Nissenbaum issued a clarion call concerning the erosion of accountability in society due to the ubiquitous delegation of consequential functions to computerized systems. Using the conceptual framing of moral blame, Nissenbaum described four types of barriers to accountability that computerization presented: 1) “many hands,” the problem of attributing moral responsibility for outcomes caused by many moral actors; 2) “bugs,” a way software developers might shrug off responsibility by suggesting software errors are unavoidable; 3) “computer as scapegoat,” shifting blame to computer systems as if they were moral actors; and 4) “ownership without liability,” a free pass to the tech industry to deny responsibility for the software they produce. We revisit these four barriers in relation to the recent ascendance of data-driven algorithmic systems — technology often folded under the heading of machine learning (ML) or artificial intelligence (AI) — to uncover the new challenges for accountability that these systems present. We then look ahead to how one might construct and justify a moral, relational framework for holding responsible parties accountable, and argue that the FAccT community is uniquely well-positioned to develop such a framework to weaken the four barriers.

Additional Key Words and Phrases: accountability, moral philosophy, robustness, data-driven algorithmic systems

1 INTRODUCTION

In 1996, writing against the backdrop of the meteoric rise of the commercial Internet [85], Helen Nissenbaum warned of the erosion of accountability due to four barriers inimical to societies increasingly reliant on computerized systems [102]. These barriers are: “**many hands**,” borrowing a term from philosophy to refer to the problem of attributing moral responsibility for outcomes caused by multiple moral actors; “**bugs**,” a way software developers might shrug off responsibility by suggesting software errors are unavoidable; “**computer as scapegoat**,” shifting blame to computers as if they were moral actors; and “**ownership without liability**,” a free pass to the software industry to deny responsibility, particularly via shrink-wrap and click-wrap Terms of Service agreements. Today, twenty-five years later, significant work has been done to address the four barriers, for example, through developments in professional practices of computer science [67, 134], organizational management [68], and civil law [43, 97]; however, on the whole, the effort to restore accountability remains incomplete. In the interim, the nature of computerized systems has once again been radically

transformed — this time by the ascendance of data-driven algorithmic systems¹ characterized by machine learning (ML) and artificial intelligence (AI), which either have replaced or complemented rule-based software systems, or have been incorporated within them as essential elements [12, 27, 82, 98, 149].

The resurgent interest in accountability, particularly among the FAccT community, is therefore timely for a world in which data-driven algorithmic systems are ubiquitous. In domains as varied as finance, criminal justice, medicine, advertising, entertainment, hiring, manufacturing, and agriculture, these systems are simultaneously treated as revolutionary, adopted in high-stakes decision software and machines [6–8, 75], and as novelties [74]. The failure to comprehensively establish accountability within computational systems through the 1990s and 2000s has therefore left contemporary societies just as vulnerable to the dissipation of accountability, with even more at stake. We remain in need of conceptual, technical, and institutional mechanisms to assess how to achieve accountability for the harmful consequences of data-driven algorithmic systems — mechanisms that address both *whom* to hold accountable and *how* to hold them accountable for the legally cognizable harms of injury, property loss, and workplace hazards, and the not-yet-legally-cognizable harms increasingly associated with data-driven algorithmic systems, such as privacy violations [30], manipulative practices [5, 80], and unfair discrimination [6]. In light of growing concerns over accountability in computing, our paper revisits Nissenbaum’s “four barriers to accountability” to assess whether insights from that work remain relevant to data-driven algorithmic systems and, at the same time, how the ascendance of such systems complicates, challenges, obscures, and demands more of technical, philosophical, and regulatory work.

We next provide necessary background concerning recent developments standards of care, law and policy, and computer science that will assist us in our analysis. Equipped with this context, we orient ourselves in a broad range of disciplines (**Section 2**) in order to examine each of Nissenbaum [102]’s four barriers to accountability in relation to data-driven algorithmic systems (**Section 3**). Having re-introduced and updated the discussion of the barriers, we then integrate our conceptual framing to suggest how one might construct and justify a moral, relational framework for holding responsible parties accountable, and argue that the FAccT community is uniquely well-positioned to develop such a framework to weaken the four barriers (**Section 4**).

Authors’ addresses: A. Feder Cooper, Cornell University, Ithaca, NY, USA, afc78@cornell.edu; Benjamin Laufer, Cornell Tech, New York, NY, USA, bd156@cornell.edu; Emanuel Moss, Cornell Tech and Data & Society Research Institute, New York, NY, USA, emanuelmoss@cornell.edu; Helen Nissenbaum, Cornell Tech, New York, NY, USA, hn288@cornell.edu.

¹Since rule-based software systems are also “algorithmic,” we will take care to specify which of the meanings we intend in settings where the context of use does not disambiguate.

1.1 Contemporary Interventions in Accountability and Data-Driven Algorithmic Systems

Re-visiting the four barriers requires engaging with the significant body of contributions concerning accountability produced in the interim. Rather than comprehensively reviewing existing literature — an immense undertaking already addressed in work such as Wieringa [145] and Kohli et al. [77] — we touch on three areas of work that we find useful for our analysis of the four barriers in Section 3: standards of care, law and policy, and computer science.

Standards of care. Standards of care play a crucial role in building a culture of accountability through establishing best practices. They provide formal guidelines for ensuring and verifying that concrete practices align with agreed-upon values, such as safety and reliability [102]. In engineering, standards of care dictate the behaviors and outputs that ought to be expected. Concerning data-driven algorithmic systems in particular, such standards of care have taken the form of model cards [94, 119], data sheets [55], annotations [13], audits [6], and frameworks [66] concerning the appropriate use of data and other artifacts, which are often developed and used in the production of AI and ML systems [22, 91]. Taken together, these standards of care enable accountability by making the intentions and expectations around such systems concrete; they provide a baseline against which one can evaluate deviations from expected behavior, and thus can be used to review and contest the legitimacy of specific uses of data-driven techniques.

In certain cases, scholars have re-framed accountability standards around harmed and vulnerable parties [93, 110]. This work — particularly that which focuses on transparency [39] and audits [135] — makes clear that standards of care and frameworks, while important for developing actionable notions of accountability, do not guarantee accountability on their own. Algorithmic impact assessments (AIAs) attempt to fill this gap [96]. They task practitioners with assessing new technologies [115] with respect to the impacts they produce [93]. IAs formalize accountability relationships in ways to systematically address and correct algorithmic harms.

Law and policy. Democratic legal literature on data-driven algorithmic systems generally concerns AI/ML-related harms and corresponding *ex post* interventions [6, 100]. For example, work on liability spans both anticipated harms related to new or forthcoming data-driven technology, including autonomous vehicles (AVs) and robotics [3, 7, 42, 128], and not-yet-legally-cognizable harms, such as unfair discrimination due to demographically-imbalanced, biased, or otherwise-discredited training data [63, 103, 140], privacy violations [30, 35, 71], and manipulation [80]. Regulatory and administrative scholarship tends to analyze data-driven algorithmic systems in relation to pre-existing legislation and policy [98, 112, 117, 137, 144]. One notable exception is GDPR — the nascent, yet wide-reaching data-privacy policy in the EU — which has also been applied to AI/ML systems [59, 72, 139].

Transparency is intimately connected to accountability. It is necessary for identifying responsible parties (in order to attribute harms to those who are responsible for them), and necessary for identifying the means by which harms came about (so that harms can be mitigated) [39]. Transparency is therefore of broad import in democratic governance [97], and therefore in law and policy. This work

spans a variety of urgent issues regarding lack of transparency in data-driven algorithmic systems. Concerning data [88, 138], lack of transparency is implicated in the obfuscation of data provenance, particularly via the concentration of data ownership within data brokers [44, 83, 148]; concerning algorithms and models, insufficient transparency contributes to the inscrutability of algorithmic decisions [31, 82, 84]. As a result, outsourcing legal decisions to automated tools, particularly data-driven tools that obscure underlying logic, can create a crisis of legitimacy in democratic decision-making [26, 29, 98, 133].

Computer science. Research in computer science, especially in ML, has increasingly treated accountability as a topic for scholarly inquiry. In updating Nissenbaum [102]’s barriers, we address cases in which researchers explicitly recognize the relationship between their work and accountability [73] — namely, in auditing and transparency — and work concerning *robustness*, which we identify as having significant implications for accountability, even when this work does not make this claim explicitly. Concerning auditing, recent work underscores the importance of being able to analyze algorithmic outputs to detect and correct for the harm of unfair discrimination [4, 110]. Transparency tends to be treated as a property of models, particularly with regard to whether a model is interpretable or explainable to relevant stakeholders [14, 40, 52].² More recently, computational work has begun to take a more expansive view of transparency, applying it to other parts of the ML pipeline, such as problem formulation, data provenance, and model selection choices [50, 50, 81, 121, 122]. Lastly, broadly construed, *robustness* concerns whether a model behaves as expected — under expected, unexpected, anomalous or adversarial conditions. Designing for and evaluating robustness implicates accountability, as it requires researchers to rigorously define their expectations concerning model performance; this in turn enables lines of inquiry aimed at both preventing deviations from those expectations, and identifying (and ideally correcting for) such deviations when they do occur. Robustness thus encompasses work in ML that aims to achieve proven theoretical guarantees in practice [69, 92, 147, 150], and work that, even in the absence of such guarantees, produces models that exhibit reproducible behavior empirically [19, 109]. Robustness also includes the ability for ML models to generalize beyond the data on which they were trained [65, 101], ranging from natural cases of distribution shift [76, 104] to handling the presence of adversaries that are trying to control or game model outputs [34, 57, 105, 129].

2 CONCEPTUAL FRAMING

We now turn to the conceptual framing we draw on for the remainder of the paper. We contextualize notions of accountability in relation to two areas: 1) contemporary moral philosophy, which situates accountability in a relationship between multiple actors; 2) the so-called “narrow” definition of accountability, largely influenced by Mark Bovens, whose framework for identifying accountability relationships has received attention from scholars interested

²It has long been contested whether metrics for model interpretability are sufficient for human comprehension [28].

in “algorithmic accountability,” especially within the FAccT community [69, 81, 145].

2.1 Accountability in Moral Philosophy

There have been various attempts in moral philosophy to develop rigorous notions of accountability. We focus on two threads in the literature — *blameworthiness* and *relationships between moral actors*, and correspondences between the two.

Blame. Nissenbaum [102] anticipated a problem of diminishing accountability as societies become increasingly dependent on computerized systems. The concern in question, simply put, was a growing incidence of harms to individuals, groups, or even societies for which no one would step forward to answer. In explaining how and why barriers to accountability emerge, the article turned to the work of legal philosopher Joel Feinberg to ground its conception of accountability. According to Feinberg, blame, defined in terms of *causation* and *faultiness*, is assigned to moral agents for harms they have caused due to faulty actions [46, 47].³

Following Feinberg, Nissenbaum [102] conceives of actors as accountable when they step forward to answer for harms for which they’re blameworthy. This conception of accountability focuses on the circumstances in which harm arises and attempts to connect an agent to the faulty actions which brought about harm. Accordingly, the barriers to accountability that Nissenbaum [102] identifies arise because the conditions of accountability are systematically obscured, due, at some times, to circumstances surrounding computerization, and at other times, to a societal breakdown in confronting willful failures to step forward. *Many hands* obscures establishing lines of causal responsibility (Section 3.1); *bugs* obscures the classification of errors as instances of faulty action (Section 3.2); *scapegoating computers* obscures answerable moral actors by misleadingly (or mistakenly) attributing moral agency to non-moral causes (Section 3.3); and *ownership without liability* bluntly severs accountability from blame (Section 3.4).

Relationality. Whereas the conception of accountability based on blameworthiness focuses on the actions and the actor causing harm through faulty behavior, recent philosophical work expands its focus to consider responsibility in light of the relationships between moral actors. Watson [142], for example, argues that responsibility should cover more than *attributability*, a property assigned to an actor for bringing about a given outcome [130]. A second dimension, which he calls *accountability*, situates responsibility in a *relationship among actors*. For Watson, “Holding people responsible

is not just a matter of the relation of an individual to her behavior; it also involves a social setting in which we demand (require) certain conduct from one another and respond adversely to another’s failures to comply with these demands” [142, p 229]. Other work, including T.M. Scanlon’s theory of responsibility, provides accounts of both *being responsible* and *being held responsible*, where the latter describes situations when parties violate relationship-defined norms [113, 120]. The moral obligations that actors have to one another are tied to the placement of responsibility *and* blame. Accordingly, the characteristics of a harmed party might dictate whether, or what, accountability is needed. For instance, if one harms in self defense, there may be no moral imperative to hold them accountable.

The works discussed attempt to situate accountability in the relationships — social, political, institutional, and interpersonal — in which we are enmeshed. Accordingly, the relationship-defined obligations we have to one another — as spouses, citizens, employees, or friends — may dictate what it is we are responsible for, as well as the contours of accountability we can expect. By situating accountability not just as *attributability* between action and actor, but instead within a social framework, some of what has come out of the so-called “narrow” notion of accountability in political theory (discussed in Section 2.2) can be derived from the vantage of a more “pure” moral philosophy. **Rather than formally pursuing this derivation here, we instead simply suggest that these notions of accountability need not be framed as alternatives to one another.** Moral philosophy, then, can offer a conceptual infrastructure through which a given relational framing — be it interpersonal, institutional, or political — can be said to be *legitimate* and ethically viable. Similarly, for practitioners holding a variety of organizational positions, the moral responsibilities that individuals hold can shape the ethical obligations and specific forms of accountability at play.

2.2 A “Narrow” Definition

In the past few years, “algorithmic accountability”⁴ has attracted growing interest in approaches that are institutional or structural in character. The work of political scientist Mark Bovens, particularly his “narrow definition” [20, 21] of accountability, has informed recent work on accountability for “algorithmic systems” [145]. Prompted by a concern that newly formed governmental structures and public authorities comprising the European Union lack “appropriate accountability regimes” [21, p. 447], Bovens proposes an accountability framework that comprises two key actors: an *accountable actor* and a *forum*. In the face of certain conditions, or in the wake of certain incidents, an accountable actor has an obligation to a forum to explain and justify itself — to address a forum’s questions and judgments, be judged by a forum and possibly suffer sanctions. Bovens calls this a “relational” definition⁵ because it defines accountability as a social relation between one actor (e.g. governmental department, a board, an agency, a committee, a public authority, or a person acting in an official capacity) and another actor

³Although neither of the two core elements is straightforward — in fact, the subjects of centuries of philosophical and legal thinking — we will reach deeper only when our core lines of inquiry demand it. Dating back thousands of years, philosophers have grappled with causation as a metaphysical challenge and a moral one of assigning causal responsibility. To this day the fascination holds, claiming the attention of computer scientists [58]. The concept of faultiness, likewise, is complex. To begin, it presumes free agency, itself a state whose metaphysical character and its role in moral attribution the subject of centuries’ long debate. Among contemporary writings on free action, Harry Frankfurt’s work [51], for example, has added further twists and expanded lines of scholarship. Faultiness is a basic concept in all legal systems and courts as well as legal scholarship have weighed in on it giving rise to widely adopted categories of harmful actions as intentional, reckless, and negligent, which have informed judgements of legal liability [45, 46, 95]. The elements of recklessness and negligence are particularly relevant to our analysis because they presume implicit or explicit standards, or expectations, that an actor has failed to meet.

⁴Concerns with the phrase will be discussed in the “Scapegoat” section of this paper.

⁵We would argue that Bovens’s *relational* framework is *social* in the sense given by Watson [142], mentioned above.

(e.g. a governmental entity, watchdog group, oversight committee, or even an individual person acting in a capacity, e.g. journalist.) The actor — accountable actor and forum — are defined in relation to each other. **Unlike in the work of the moral philosophers discussed above (Section 2.1), Bovens’s framework is *not* directed at rights and obligations we have to one another as moral actors, but rather as actors in roles and capacities defined by the respective sociopolitical structures in which we live.**

Bovens’s approach to accountability has emerged as an attractive framework for promoting accountability in societies that have given over key controls and decisions to AI/ML, as it has been recognized in FAccT literature that a relational framework may compel clarity on murky issues [69, 145]. For one, as discussed in Wieringa [145], Bovens directly illuminates the sociopolitical stakes of transparency, explainability, and interpretability, illustrating why these concepts are necessary for any accountability framework for data-driven societies, even though they are ultimately not sufficient to constitute accountability in and of themselves.⁶ For another, Bovens allows us to highlight parameters for what constitutes accountability, and for which appropriate values need to be specified: i.e., *who* is accountable, *for what*, *to whom*, and *under which circumstances*? The values for these parameters may be deeply contextual, and the work of tuning these parameters may, as suggested by Metcalf et al. [93], lie in the sociopolitical contestations, the “slow boring of hard boards” [143, p. 225], by the many constituencies implicated in any particular computational system.

The domain of philosophical work that could inform a moral conception of accountability is vast but not decisive in its nuances. Accordingly, we hold onto the bare bones of Joel Feinberg’s notion of blameworthiness (Section 2.1), with one amendment, as we revisit the four barriers in light of data-driven algorithmic systems. **The amendment shines a spotlight on moral claims of the affected, or harmed parties whose position in relation to those whose actions have harmed them may obligate the latter to account for this.** This work in moral philosophy aligns with work on accountability as a property of social structures, which holds it to be relational — not merely as a requirement on an accountable party to “own up” to blameworthy action as an obligation *to* another. We reserve a fuller discussion of work on accountability as a property social arrangements for Section 4. **For now, we merely remark that this work is not an *alternative* to the conception of accountability as stepping forward for blameworthy action. Rather, the two are co-existent, each applying to a different object: one an attribute of moral actors, the other, an attribute of institutional, or societal structures.**

As we demonstrate below in Section 3, we find that data-driven algorithmic systems heighten the barriers to accountability as establishing the conditions of blame — causal responsibility and fault — are even further obscured and, in turn, the ability for blameworthy actors to obfuscate their roles reduces pressure to step forward.

⁶We return to this idea in Section 4.

3 REVISITING THE FOUR BARRIERS TO ACCOUNTABILITY

In a typical scenario in which software is integrated into a functional system — fully or partially displacing (groups of) human actors — it may not be obvious that accountability could be displaced along with human actors who, frequently, are its bearers. The cumulative effect of such displacements is the increasing incidence of harmful outcomes for which no one answers, whether these outcomes are major or minor, immediate or long-term, or accrue to individuals or to societies. Resuscitating accountability is no simple task — so Nissenbaum [102] argues — because computerization sets up particularly troublesome barriers to accountability: **Many hands** (3.1), **Bugs** (3.2), **The computer as scapegoat** (3.3), and **Ownership without liability** (3.4). These interdependent barriers are not necessarily an essential quality of computer software. Rather, they are a consequence of how software is produced, integrated into institutions, and embedded within cyber-physical systems; they are a function of the wonderment and mystique that has grown around computerization, and the prevailing political-economy within which the computer and information industries have thrived. In the sections that follow, we revisit these barriers to accountability with an eye turned toward their implications amidst the massive growth and adoption of data-driven algorithmic technologies. We provide examples of the barriers in action, and defer discussion of how the barriers can be weakened to Section 4.

3.1 The Problem of Many Hands

The barrier of many hands concerns the issue that many actors are involved in the design, development, and deployment of complex computerized systems. When these systems cause harm, it can be difficult to isolate the component(s) at the source of that harm. In the absence of an accountability framework, this difficulty can conceal which specific individuals should step forward to take responsibility for the harm. As Nissenbaum summarized, “Where a mishap is the work of ‘many hands,’ it may not be obvious who is to blame because frequently its most salient and immediate causal antecedents do not converge with its locus of decision making” [102, p. 29].

Nissenbaum further analyzes the difficulty of the barrier of many hands by showing how it operates at four different levels: 1) software is produced in institutional, often corporate, settings in which there is no actor responsible for all development decisions; 2) within these settings, multiple, diffuse, groups of engineers contribute to different segments or modules of the overall deployed system, which additionally often depends on software designed and built by other actors (which, in today’s landscape, may result in licensed or freely-available open-source software); 3) individual software systems often interact with or depend on other software systems, which themselves may be unreliable or present interoperability issues; 4) hardware, not just software, often contributes to overall system function, particularly in cyber-physical systems, and it can be difficult to pinpoint if harms occur due to issues with the code, the physical machine, or the interface between the two. Any and all of these four levels of many hands problems can operate simultaneously, further obscuring the source of blame.

These particular difficulties of the problem of many hands remain today; moreover, they have been further complicated in numerous ways, given that computer systems are now ubiquitous rather than ascendant. We focus our discussion on how data-driven algorithm systems further complicate this barrier, using two illustrative (though necessarily non-exhaustive) examples. First, we discuss how the *machine learning pipeline* — the multi-stage process by which machine-learned models are designed, trained, evaluated, and deployed — presents novel problems concerning many hands. Second, contemporary data-driven algorithmic systems rely significantly on the composability of openly-available ML toolkits and benchmarking suites; these toolkits, often developed and maintained by large tech companies, tend to be advertised as general- or multi-purpose, and are frequently (mis)used in specific, narrow applications.

The machine learning pipeline. The ML pipeline is a dynamic series of steps — each of which can involve multiple (groups of) actors, including designers, engineers, managers, researchers, and data scientists. The pipeline typically starts with problem formulation and, in commercial settings, results in the deployment (and continued monitoring) of a trained model [106]. Problem formulation can involve the collection, selection, or curation of a dataset, and then the operationalization of a concrete task to learn, such as classifying loan-granting decisions or generating natural-language text. The actors responsible for formulation may then hand off their work to others responsible for implementation: choosing the type of model to fit to the data and the learning procedure to use for model training. In selecting the type of model to train, these actors may custom-design their own architecture, or may defer to using a pre-existing model, such as an off-the-shelf neural network, which has been designed by others, potentially at another company or institution. The stage of model training and evaluation can then begin, in which a set of actors runs training multiple times, perhaps for multiple combinations of chosen model types, training procedures, and hyperparameter values [121]. These actors compare trained models, from which they select some “best”-performing model (or ensemble of models), where “best” is often informed by a chosen quantitative metric, such as mean overall accuracy [50]. These stages of the ML pipeline, from formulation to evaluation, are often repeated in a dynamic loop: Until the model passes the threshold of some specified performance criteria, the process can cycle from (re)-modeling to tuning and debugging. Finally, if the model is deployed in practice, yet another set of actors who monitor the model’s behavior in practice, ensuring that it behaves in a way that aligns with expectations developed during training and evaluation.

Each stage of the ML pipeline involves numerous actors — in fact, potentially *uncountably* many actors if the pipeline employs third-party model architectures or ML toolkits, which we discuss below.⁷ Thus, if a trained model causes harms in practice, it can be extremely challenging to tease out particular actors who should answer for them. For example, harms could originate from how actors operationalize the learning task at the start of the pipeline [32],

⁷Participatory design further expands the set of many hands to end-user stakeholder [122]. This illustrates an additional manifestation of the barrier: when harms occur, it is possible to shift blame to harmed end-users, as they were explicitly involved in the ML pipeline.

move from high-level abstraction to concrete implementation [116], select hyperparameters or random seeds during model selection [50]. Accountability could lie with actors in any part of the pipeline, or some combination thereof. Bias could creep in early on from the choice of dataset, and then accumulate and become magnified further downstream during model selection. In other words, the diffuse and dynamic nature of the pipeline makes locating accountability extremely challenging. This can be understood as an issue of transparency — beyond the specific the problem of model interpretability — concerning *who* is responsible *for what*, and *how* this can be related to overarching accountability with respect to a model’s ultimate use in practice [81].⁸

Multi-purpose toolkits. Practitioners and researchers often do not code model architectures or optimization algorithms from scratch. Just like how Nissenbaum identified a problem of many hands due to the integration of third-party software modules, builders of data-driven algorithmic systems today often rely on toolkits produced by others. Such toolkits, which demand tremendous mathematical and software engineering expertise to develop, are arguably indispensable to individual and small-business model developers and data scientists. To decrease the amount of time and money spent iterating the ML pipeline, these actors depend on the investment of large tech companies with vast resources and large, concentrated pools of technical talent to develop and release efficient, correct, comprehensive, and user-friendly libraries of algorithm implementation, model architectures, and benchmark datasets [1, 107]. Of particular note is that, due to the time and expense of training increasingly larger models, some of these toolkits contain pre-trained models — large-language models like BERT [38], which can be used out-of-the-box or fine-tuned for particular use cases. Since such models are conceived of as multi- or general-purpose, they have been shown to exhibit bias-related harms when used in narrow application contexts [78, 89, 99]. Determining blame for these types of harms is far from simple; for example, if intended use is under-specified, blame could lie at least partially with the model creator.

3.2 Bugs

Nissenbaum [102] uses the term “bug” to cover a variety of issues common to software, including “modeling, design, and coding errors.” Bugs are said to be “inevitable,” “pervasive,” and “endemic to programming” — “natural hazards of any substantial system” [102, p. 32]. Even with software debuggers and verification tools that can assure correctness, “bugs” emerge and cause unpredictable behavior when software systems are deployed and integrated with each other in the real world. [90, 124]. In short, “bugs” are predictable in their unpredictability; they serve as a barrier to accountability because they cannot be helped (except in obvious cases), and therefore are often treated as an accepted “consequence of a glorious technology for which we hold no one accountable” [102, p. 34]. As Nissenbaum [102] notes, what we consider to be the “inevitable” can change over time as technology evolves, with certain types of “bugs” spilling over into the avoidable. For example, evolving norms

⁸This indicates why transparency in the form of model interpretability may be important, but is ultimately not sufficient, for identifying actors accountable for harms; irrespective of interpretability, the pipeline can muddle transparency.

and new debugging tools can rebrand the “inevitable” to be sloppy or negligent implementation, at which point programmers can be held to account for such errors. Similarly, the advent of data-driven algorithmic systems has recently indicated that this malleability also extends in the other direction: New technological capabilities can both contract and expand what we consider “inevitable” buggy behavior. That is, while these systems contain “bugs” of the “modeling, design, and coding” varieties that Nissenbaum [102] describes for rule-based programs, the statistical nature of data-driven systems presents additional types of harm-inducing errors, which may present an additional barrier to accountability.⁹ For example, misclassifications, statistical error, and nondeterministic outputs may cause harm, and likewise may be treated as inevitable, which makes it difficult to place blame.

In relation to the treatment of “bugs” in 1996, it is important to note that labeling these types of errors “bugs” presents a complication, as they are an inherent part of machine learning attributable to its statistical nature. That is, misclassification, statistical error, and nondeterminism seem to turn the notion of “bug” on its head: Many experts would just as readily call these *features* of machine learning, rather than *bugs*.¹⁰ Nevertheless, regardless of where one attempts to draw the line, these errors share common elements with the “bugs” Nissenbaum [102] describes — namely, they fit into an overarching category of software issue for which the ability to reason about causality and fault is elusive. In other words, they present a barrier to accountability due to being treated as “inevitable,” “pervasive,” and a “consequence of a glorious technology for which we hold no one accountable” [102]. To clarify this barrier, we next consider some concrete examples of “bugs” that data-driven algorithmic systems present.

Faulty modeling premises. Prior to implementation, as discussed in Section 3.1, data-driven algorithmic systems require significant modeling decisions. For example, choosing a model to learn necessarily involves abstraction and can have significant ramifications [106, 116]; assumptions during this stage of the ML pipeline can bias the resulting computational solution space in a particular direction [53]. For example, assuming a linear model is sufficient to capture patterns in data precludes the possibility of modeling nonlinearities. When such biases involve over-simplified or faulty reasoning, they can result in model mis-specification and the introduction of “modeling error bugs.” Such mis-specifications may include the assumption that values like fairness and accuracy are correctly modeled as a trade-off to be optimized [32], and that physical characteristics can serve as legitimate classification signals for identifying criminals [146] or inferring sexual orientation [126, 141]. More generally, a common modeling error involves the assumption that a problem is amenable to classification — that it is possible to divide data examples into separable categories in the first place [123, 127]. Being grounded in such false premises means that, even if it is possible to train mis-specified models like these to behave “accurately” (i.e., to return better-than-chance results after learning these tasks), the conclusions we can draw from them are unsound [32]. In these

⁹Of course, statistical software is not new to ML; however, the proliferation of data-driven algorithmic systems has clarified the prevalence of such errors.

¹⁰We return to this idea in Section 3.3, in which we discuss the accountability barrier of treating properties of ML as a scapegoat.

cases, if modeling assumptions are unclear or elided, such that one is not able to attribute errors to them, it is easy to evade accountability by placing blame on the presence of inexplicable, unavoidable “bugs” endemic to computer software.

Individual errors. Avoiding faulty premises is not alone sufficient to guarantee that the ML pipeline produces a harm-free model, as such models tend to exhibit some level of error leading to individual instances of harm of different varieties, including disparate impact or manipulation [48, 80, 89, 99]. ML has several metrics to quantify error [17, 61]; in training, an optimization algorithm attempts to minimize a chosen error metric. Nevertheless, even the most robust, well-trained models report imperfect accuracy. In fact, a model that achieves 100% accuracy is usually considered suspect; it likely overfit to the training data and to exhibit poor performance when presented with new examples [62, 125]. Therefore, when individual errors occur, they can be treated as inevitable, just like the “bugs” Nissenbaum [102] describes, thus displacing responsibility for the harm such errors cause affected individuals.

Bad model performance. Like the problem of excusing individual errors as unavoidable “bugs,” it is possible to treat unexpectedly bad overall model performance similarly. Consider a hypothetical example of a (well-formulated) computer vision system used to detect skin cancer, whose training and evaluation promise will have a mean accuracy rate of 94%. If, once deployed, the model coheres with (or out-performs) its promised performance, then misclassifications can be said to have been anticipated or *expected*.¹¹ Since expected accuracy is a probabilistic claim about what is likely to occur, deviations from expectation can (and do) occur. If, when monitoring a deployed model, this deviation yields a decrease in expected model accuracy that is sustained over time, it is possible to evade accountability by ascribing the issue to the amorphous category of “bug.” That is, calling this type of error an inevitable “bug” avoids attributing under-performance to a particular source, which, rather than being unavoidable, could be a result from human negligence, poor generalization, distribution shift, or other faulty behavior.

3.3 The Computer as Scapegoat

Nissenbaum [102] suggested that the practice of blaming a computer could pose a barrier to accountability, since “having found one explanation for an error or injury, the further role and responsibility of human agents tend to be underestimated” [102, p. 34]. To explain why people could plausibly blame computers for a wrongdoing, Nissenbaum points to the fact that “computers perform tasks previously performed by humans in positions of responsibility”; whereas before the human would be indicated as the blameworthy party, the computer has now taken up the role. And yet, while computer systems have become more immediate causal antecedents to an expanding number of harms, they do not have moral agency and thus cannot be said to conduct ethically *faulty* behavior [102]. We discuss how the barrier of scapegoating the computer has become

¹¹Individual instances of error can pose additional challenges for accountability, since the model may still overall exhibit an expected degree of error (i.e., be within a margin of error); in these cases, it is possible to treat “the computer as scapegoat” — another barrier to accountability — and say error is inherent to the fundamental statistical nature of ML. We discuss this further in Section 3.3.

more complicated within the landscape of ubiquitous data-driven algorithmic systems: the tendency to treat such systems as “intelligent” objects with moral agency (which they do not possess) provides unprecedented opportunities to scapegoat the system or its component algorithms [108, 132]. In seeming opposition to treating the system as intelligent or rational, the non-determinism and stochasticity inherent in ML can cause inexplicable behavior or errors, which also presents an occasion to misplace blame.

Moral agency. As data-driven algorithmic systems have become pervasive in life-critical contexts, there has been a corresponding tendency to anthropomorphize and equate technological processes to human cognition [108, 132]. The “intelligence” in artificial intelligence and the “learning” in machine learning would suggest a sort of adaptive, informed decision apparatus that enables a neat placement of moral blame. However, directing blame toward data-driven algorithmic systems effectively imbues them with moral agency, ascribing them the ability to act intentionally [114].¹² Nissenbaum [102] likens blaming a computer to blaming a bullet in a shooting: While the bullet can be said to play an active, causal role, it cannot be said to have been intentional in its behavior. In the same vein, a data-driven algorithmic system may play a central role in life-critical decisions, and may even be said to *make a choice* in a particular task, but such a choice does not hold a deliberate intention that constitutes moral agency [114].¹³

“Accountable algorithms”. This banner-phrase makes algorithms the *subject* of accountability, even though algorithms cannot be said to hold moral agency and, by extension, moral responsibility. This popular term [82], therefore, reduces accountability to a piecemeal, procedural quality that can be deduced from a technology, rather than a normative concept that has to do with the moral obligations that people have toward one another. Moreover, algorithms do not get deployed in practice; systems (which contain algorithms) do. When, for example, studies fairness in AI/ML-assisted judicial bail decisions fixate on the biases that exist *within* an algorithm, they fail to capture broader inequities that are systemic in complex sociotechnical systems, of which AI/ML techniques are just one part [2, 10, 33].

Mathematical guarantees. This attempt to direct blame away from people and corporations can be either strategic or unconscious. In some cases, a group of harmed individuals do not know who to blame (Section 3.1) and settle on blaming the system. In others, scapegoating the system can be a way by which an actor dodges and dissipates public ire. For example, consider the now-canonical example of Northpointe exhibiting bias in their risk-assessment

tool [8]; rather than attributing this bias to a mistake or “bug,” Northpointe blamed the fundamental incompatibility of different algorithmic operationalizations of fairness as the source of the problem (and pointed to a specific measure, for which bias was not detectable, as evidence of blamelessness).

More generally, mathematical guarantees, if blamed for harmful outcomes, can exhibit the barrier of scapegoating. In contrast to the example of unexpectedly poor model performance, described above in the “Bugs” (Section 3.2, consider the following case: Engineers design a data-driven algorithmic system which they analytically prove — and empirically validate — meets some specified theoretical guarantee. In particular, let us even consider the same case we discussed for the problem of individual errors in “Bugs”: The engineers prove that a system is 94% accurate in detecting tumors, and then validate that this is in fact the case in practice. This same example, depending on how it is unpacked, can exhibit the barrier of “bugs” or the barrier of scapegoating the computer. Above, we talked about this example in terms of individual errors, for which responsibility for harm could be excused due to “buggy” behavior. Here, rather than analyzing behavior at the level of individual decisions, we examine the behavior *of the model overall*. If the frequency of mis-classifications is within the model’s guaranteed error rate, the engineers could attempt to excuse all resulting harms by gesturing to the fact that the model is performing *exactly as expected*. In short, satisfying mathematical guarantees can serve as a scapegoat because pointing to mathematical claims satisfied at the model-level can serve to obscure the need to account for harms that occur at the individual-decision level.¹⁴

Non-determinism. Data-driven algorithmic systems that involve ML exhibit non-determinism. They involve randomization to, for example, shuffle the order in which training data examples are presented to an algorithm. While such features of ML algorithms may seem like technical minutiae, they in fact introduce stochasticity into the outputs of machine-learned models: Training the same model architecture on the same dataset with the same algorithm — but changing the order in which the training data are supplied to the algorithm — can yield models that behave very differently in practice [37]. For example, as Forde et al. [50] shows, changing the order that the data examples are presented to a train tumor-detection model can lead to surprisingly variable performance. The relationship between training-data-ordering and the resulting variance in model performance is under-explored in the technical literature. Thus, such differences in model performance are often attributed to an inherent stochasticity in ML. The randomization used in ML algorithmic systems — randomization on which these systems depend — becomes the scapegoat for the harms it may cause, such as missed tumor detection.

¹²Proponents of the standard conception of agency in Schlosser [114] include Davidson [36], Goldman [56], and Brand [23]. It is also applied to artificial agency by Himma [64].

¹³This is consistent with scholarship in both moral philosophy and legal theory concerning AI, algorithms, agency, and personhood [15, 24, 64, 136]. The legal literature has called this, particularly in relation to robots, “social valence” [25] and “The Substitution Effect” [9]. Moreover, since in the US tort law is grounded in a notion of moral culpability, it has been ill-suited for application to AI-harm-related remedies (see, e.g., [86]).

¹⁴And, of course, one can see-saw back-and-forth between “bug” and scapegoating to evade accountability. If satisfying guarantees at the overall model-level is for some reason rejected as a rationale for an individual harm, one could claim there is a “bug”; if calling an individual decision “buggy” is rejected, and the model is classifying within an expected threshold for error, one could then displace blame by arguing that the model is according to its specification.

3.4 Ownership without Liability

Nissenbaum [102] highlights a dual trend in the computer industry: 1) strengthening property rights; 2) avoiding liability. Behavioral trends that informed these assertions have persisted in the decades since, with lively public debates over the fit of traditional forms of intellectual property (i.e., copyrights, patents, and trade secrets) to digital products such as software, data, databases, and algorithms, and subsequent expensive legal struggles among industry titans [16]. Similarly, we have seen explicit denials of liability expressed in shrink-wrap licenses, carried over into so-called “click-wrap” licenses, and terms of service disclaimers accompanying websites, web-based services, mobile apps, IoT devices, content moderation decisions, and the like [30, 79, 87, 131].

Before addressing how we see these trends carrying forward in the contemporary landscape, we need to qualify our observations. Property and liability are weighty legal concepts with long histories and rich meanings. Narrowing our view to digital technologies, even before the 1996 paper a robust literature had grown over questions of ownership — questions that have persisted through numerous landmark court cases. Liability, too, is a core legal concept that is increasingly an issue in relation to the products and services of digital industries. It lies outside the scope of this paper to attempt meaningful insights into these concepts as they manifest in scholarship, law, and the courts. However, it is useful to observe broad patterns and anticipate the likely actions of stakeholders.

For a start it is not difficult to see how the trends toward strong ownership and weak liability reinforce barriers to accountability, and also to understand why industry incumbents might support them; liability is costly and strong property rights enrich rights holders and empower them against competitors. Four lines of advocacy on behalf of industry interests are noted below, supplementary to those discussed in Nissenbaum [102]:

- (1) Third-party providers of data-driven algorithmic systems refuse to expose their systems to scrutiny by independent auditors on grounds of trade-secrets [31, 54]. As long as experts maintain that transparency is necessary to evaluate the ML pipeline and AI development, strong property rights that block scrutiny are barriers to accountability.
- (2) Manufacturers and owners of cyber-physical systems, such as robots, IoT devices, drones, and autonomous vehicles, evade liability for harms by shifting blame to environmental factors or humans-in-the-loop [86]. In this respect, the barrier of ownership without liability for data-driven algorithmic systems suggests a twist on the problem of scapegoating (Section 3.3): treating “the human user as scapegoat.” That is, claiming the user has mis-used an AI- or ML-enabled system in order to obscure responsibility for unclear, under-specified, or deliberately misleading user interfaces or expected use, as has happened with Tesla and accidents concerning its (so-called) “AutoPilot” autonomous driving feature [18].
- (3) Almost without question the computer industry, having metamorphosed into the data industry, has assumed ownership over data passing through its servers [44, 83, 103]. We still do not have clear rules of liability for industry actors when their servers, holding unimaginable quantities of data, are breached [118].

- (4) Technology companies hold unprecedented sway over regulation. Twenty-five years ago, although the software industry was already a force to be reckoned with, it successfully persuaded Congress that imposing legal constraints would stifle innovation — that societal well-being depended on a nascent industry that could not flourish under excessive regulatory and legal burden. Despite the obvious maturing of the industry and the emergence of global industry titans the innovation argument seems to hold sway [111], this time (in the US) with a twist: concern over losing tech market dominance to emerging economies.

4 BEYOND THE FOUR BARRIERS: LOOKING AHEAD

Nissenbaum [102] warned of a waning culture of accountability — harms befalling individuals, groups, even societies, accepted as sufferers’ bad luck, because no one would be stepping forward to answer for them. In the previous section, we revisited the four barriers in light of data-driven algorithmic systems and found that the framework still provides a useful lens through which to locate sources of the dissipation of accountability. Striking down, or even weakening the barriers would clear the way for sound attribution of blame, in turn exposing blameworthy parties to a societal expectation to step forward. But we have also argued that the moral responsibility to step forward is a necessary, but insufficient, component of accountability. In spite of, or perhaps because of the barriers, established institutional frameworks calling some actors (individuals or groups) to account and empowering others to call to account, as an alternative to case-by-case attention, is a more promising approach to building a culture of accountability (Section 4.1), and demonstrate how the FAccT community is uniquely disposed to develop future work that would erode the barriers to both forms of accountability (Section 4.2).

4.1 Bringing together moral and relational accountability

A robust accountability framework needs to specify *who* is accountable, *for what*, *to whom*, and *under which circumstances*. Bovens’ notion of relational accountability suggests such a framework, including two key actors: a *forum* and an *accountable actor*, who has an obligation to justify itself to the members of the forum (Section 2.2). We have argued for the appropriateness and urgency of unifying this relational definition with a conception of blameworthiness, and now extend this argument to highlight the need for a moral, relational accountability framework for dealing with the harms of data-driven algorithmic systems. In light of this call, the moral conception of accountability, on which Nissenbaum [102] depended, suggests one partial answer: Those who have caused harm through faulty action are contenders for the class of *accountable actors*, and those who have suffered harm deserve a place among the members of the *forum*. This point usefully clarifies a confluence between accountability as answerability for blameworthy action, and accountability as a social arrangement. Being blameworthy for harm is (almost always) a sufficient condition for being designated an accountable actor; being harmed through blameworthy action (almost always) is a sufficient condition for being designated a member of the forum. **These two conceptions do not stand**

against one another as alternative solutions to the same problem; they are solutions to different problems that intersect in instructive ways. We thus show how a moral, relational accountability structure thus widens the scope — beyond the pair harmed party and faulty actor — by providing examples of values for all four parameters:

Who is accountable. *Accountable actors* may include those who are not directly responsible for harm (e.g. engineers) but are designated as accountable (or liable) because of their deep pockets, capacities to render explanations, or positions in organizational hierarchies, such as corporate officers or government procurers of data-driven systems.

For what. Beyond legally-cognizable harms considered in tort law, such as bodily injury, property damage, pecuniary losses, and even non-physical harms, such as those to reputations, harms particularly associated with data-driven algorithmic systems include privacy violations [30], unfair discrimination [6], and autonomy losses due to manipulation [5, 80]. We take as an example privacy harms from AI/ML. Privacy rights of data subjects, in the creation of training and testing data sets has rightfully drawn attention from advocates of data sheet [22, 55], for example, and likewise a rationale for differential privacy [41], whose focus is the data subject. But this limited view of harmed parties fails to consider the privacy of parties affected by the uses of models, even if creation of the data from which the models are derived has followed recommended, privacy preserving standards, e.g. anonymization, contextual integrity, permission, etc. As discussed in Hanley et al. [60], a face dataset that has been created with utmost attention to privacy, may nevertheless cause privacy harm when a model derived from it is used to identify individuals as gay [141].¹⁵ In other words, once there is a trained model, it can be used to learn things about specific people; even if one were to anonymize a dataset for training, that dataset — by being used for training — can be operationalized and used to learn things about other people who are not in the dataset.

To whom. The members of the *forum* may not just include those who are themselves harmed (or placed in harms way through heightened risk). They may also include those deputized to represent and advocate on behalf of vulnerable parties, such as lawyers, public or special interest advocacy groups; or, beyond direct advocates, groups and individuals in oversight capacities, such as journalists, elected officials, government agencies, or professional societies.

Under which circumstances. This concerns the nature of the obligation, or, what *accountable actors* may owe to the forum — to explain, be judged, and address questions and challenges. Not every algorithmic system used in every domain requires the same approach to accountability. Rather, the specific responsibilities an actor has to others are inflected by the context in which they act.

4.2 Weakening the barriers

In the previous section, we laid out what would be needed to satisfy a moral and relational accountability framework. Any technical interventions that have already been developed — notably, those that

we have emphasized concerning transparency, audits, and robustness — would need to be folded into such a framework, their use justified in these moral and relational terms. For example, whatever is proposed as a technical definition of transparency, it is unlikely to satisfy the needs of all those who comprise a forum, who may not be educated in the particulars of what it means for a model to be “interpretable.” Robustness says what expectations are, but leaves unanswered the question of the conditions under which deviations from expectations ought to be remedied. Relational treatments of these issues, it would seem, require that the obligation be tuned to the variable needs of all members of the forum.

Aside from these justifying pre-existing interventions, as we have demonstrated in Section 3, new interventions are also needed to weaken the barriers to accountability. For example, a moral and relational accountability framework opens the aperture, in principle, to many, if not all, of the “many hands” being designated as accountable actors, including dataset creators, model developers, decision and control systems designers, purveyors, and operators of these systems (Section 3.1. Developing rigorous standards of care could help mitigate the problems of inappropriate use of pre-trained models and unclear measures of quality control at different stages of the ML pipeline. For example, robust auditing mechanisms at each stage, rather than treating auditing as an end-to-end concern [110], could help clarify the relationship between stage-specific issues and resulting harms.

Various harms, depending on how they are contextualized, can implicate either the barrier of “bugs” or “scapegoating the computer” (Sections 3.2 & 3.3). For example, we note that the computer science community could have treated harms due to unfair discrimination as either a “bug” or blamed them on intrinsic aspects of AI/ML — and yet they did not. In relation to “bugs,” in the past, unfairness often was ascribed due to biased or imbalanced training data [49, 70]. Such biased historical data is arguably “pervasive” and unavoidable. Similarly, the community could have set some “tolerable” level of model unfairness and, as long as model met that specification, they could have attempted to evade accountability by blaming inherent properties of the model. And yet, the community does not redirect blame for the harms of unfair discrimination to these barriers.¹⁶ While unfair discrimination remains a serious issue in data-driven systems, FAccT and its antecedents have made a significant effort not to evade accountability for unfairness harms by attributing them to “bugs” or treating the computer as scapegoat. For example, significant attention has been paid to mitigating unfairness harms by developing training algorithms that are robust to such biased input data. The field of algorithmic fairness therefore serves an example that challenges the narrative these barriers — an example that could encourage similar treatment of other issues like robustness and its relationship to privacy violations, or adversarial ML and its relationship to manipulation. The community has demanded more from ML modelers concerning the treatment of unfair discrimination; it has set expectations concerning the necessity of interventions to root out and correct for unfairness, thereby surpassing the barriers of being attributed to “bugs” or scapegoated.

¹⁵A related point is discussed in Barocas and Nissenbaum [11].

¹⁶Arguably, it would be repugnant to do so; the particularly ugly nature of unfair-discrimination-related harms may be the reason they have escaped such treatment, which could be perceived as flippant, discriminatory, and thus a harm in itself.

Lastly, being liable is related but not identical to being accountable (Section 3.4). The latter is applied to blameworthy parties who step forward to answer, the former to parties who step forward to compensate victims of harm. Often liability is assigned to those who are found to be blameworthy. If lines of accountability are blurred, for example, as a consequence of the barriers we have discussed, harms due to AI/ML and other data-driven algorithmic systems will be viewed as unfortunate accidents; the cost of “bad luck” will settle on victims. Instead, legal systems have developed approaches, such as strict liability, to compensate victims harmed in certain types of incidents even without a showing of faulty behavior. Strict liability assigned to actors who are best positioned to prevent harm is sound policy as it is likely to motivate these actors to take extraordinary care with their products. If, indeed, barriers such as many hands make the attribution of blame impossible, strict liability for a range of algorithm driven harms, such as privacy breaches, unfair discrimination, manipulative practices, as well as traditional injuries would, at least, shift the “bad luck” from victims to those best positioned to take extraordinary steps to mitigate and prevent such harms.

FACCT has a key role in developing these and other tools to help erode the barriers to accountability. However, the use of these tools needs to be justified. Just as mature political governance requires durable institutions and formal attributions of rights and duties, we have similar needs for the governance of producers, purveyors, and operators of data-driven algorithmic systems. That is, as we have contended throughout this paper, accountability is moral and relational. It depends on social, legal, and political structures that provide legitimacy for the checks actors and forums place on each others behavior; it depends on the way those checks are internalized as professional, personal, legal, and ethical duties that motivate actors’ personal responsibility. FACCT, given its proclaimed valuing of accountability and the array of expertise it brings together, is uniquely positioned to help develop a relational and moral accountability framework — the structures that provide legitimacy, as well as the professional codes and standards of care, disciplinary norms, and personal mores that tie it all together. The future work of creating these structures, as we noted earlier, is no small undertaking; it lies in the sociopolitical contestations, the “slow boring of hard boards” [143, p. 225], by the many constituencies implicated in any particular computational system.

5 CONCLUSION

In this paper we revisited Nissenbaum [102]’s “four barriers” to accountability, with attention to the contemporary moment in which data-driven algorithmic systems have become ubiquitous in consequential decision-making contexts. We draw on conceptual framing from Nissenbaum [102]’s use of the concept of *blameworthiness*¹⁷ and how it can be aligned with, rather than cast in opposition to, Mark Bovens’ work on accountability as a *relational property of social structures* [20, 21]. We demonstrate how data-driven algorithmic systems heighten the barriers to accountability with regard to determining the conditions of blame, and look ahead to how one might endeavor to weaken the barriers. In particular, drawing on

both Nissenbaum and Bovens, we put forward the conditions necessary to satisfy a moral and relational accountability framework, discuss how the development of such a framework would weaken the barriers, and argue that the FACCT community is uniquely positioned to construct such a framework and to develop lines of inquiry to erode the barriers to accountability. Given our tender historical moment, addressing why these or those parties belong in the forum or in the set of accountable actors, why those obligations are justified, and, of course, evaluating the numerous permutations the relational nature of the approach demands is the provenance of future work. No easy formulations or operationalizations make sense until we have developed a rigorous approach to justification. In our view, this calls for expertise in relevant technologies, moral philosophy, the prevailing political economy of data and computing industries, organizational sociology, prevailing political and regulatory contexts, domain area expertise, and more — areas that FACCT has successfully brought together under its sponsorship. It is not that all these are needed all the time; but any of them may be called in to develop linkages between proposed values and social welfare.

ACKNOWLEDGMENTS

The authors would like to thank the Digital Life Initiative at Cornell Tech for its generous support. A. Feder Cooper is additionally supported by the Artificial Intelligence Policy and Practice initiative at Cornell University and the John D. and Catherine T. MacArthur Foundation. Benjamin Laufer is additionally supported by NSF Grant CNS-1704527. Helen Nissenbaum is additionally supported by NSF Grants CNS-1704527 and CNS-1801307, and ICSI Grant H98230-18-D-006.

REFERENCES

- [1] Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <https://www.tensorflow.org/> Software available from tensorflow.org.
- [2] Rediet Abebe, Solon Barocas, Jon Kleinberg, Karen Levy, Manish Raghavan, and David G Robinson. 2020. Roles for computing in social change. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. ACM, New York, NY, USA, 252–260.
- [3] Kenneth S. Abraham and Robert L. Rabin. 2019. Automated Vehicles and Manufacturer Responsibility for Accidents: A New Legal Regime for a New Era. *Virginia Law Review* 105 (2019), 127–171. Issue 127.
- [4] Philip Adler, Casey Falk, Sorelle A. Friedler, Tionney Nix, Gabriel Rybeck, Carlos Scheidegger, Brandon Smith, and Suresh Venkatasubramanian. 2018. Auditing black-box models for indirect influence. *Knowledge and Information Systems* 54 (2018), 95–122.
- [5] Shruti Agarwal, Hany Farid, Yuming Gu, Mingming He, Koki Nagano, and Hao Li. 2019. Protecting World Leaders Against Deep Fakes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. IEEE, Long Beach, CA, 8.
- [6] Ifeoma Ajunwa. 2021. An Auditing Imperative for Automated Hiring. *Harv. J.L. & Tech.* 34 (2021), 81 pages. Issue 1.
- [7] American Association for Justice. 2017. Driven to Safety: Robot Cars and the Future of Liability. , 50 pages.
- [8] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine bias. *ProPublica* 23, 2016 (May 2016), 139–159.
- [9] Jack M. Balkin. 2015. The Path of Robotics Law. *California Law Review Circuit* 6 (2015), 45–60.

¹⁷ Nissenbaum [102], in turn, drew on the work of Joel Feinberg [46, 47].

- [10] Chelsea Barabas, Colin Doyle, JB Rubinovitz, and Karthik Dinakar. 2020. Studying up: reorienting the study of algorithmic fairness around issues of power. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. ACM, New York, NY, USA, 167–176.
- [11] Solon Barocas and Helen Nissenbaum. 2014. Big Data’s End Run around Anonymity and Consent. In *Privacy, Big Data, and the Public Good: Frameworks for Engagement*, Julia Lane, Victoria Stodden, Stefan Bender, and Helen Nissenbaum (Eds.). Cambridge University Press, Cambridge, UK, 44–75.
- [12] David Barstow. 1988. Artificial Intelligence and Software Engineering. In *Exploring Artificial Intelligence*. Elsevier, 641–670.
- [13] Elena Beretta, Antonio Vetrò, Bruno Lepri, and Juan Carlos De Martin. 2021. Detecting discriminatory risk through data annotation based on Bayesian inferences. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM, New York, NY, USA, 794–804.
- [14] Umang Bhatt, Javier Antorán, Yunfeng Zhang, Q. Vera Liao, Prasanna Sattigeri, Riccardo Fogliato, Gabrielle Melançon, Ranganath Krishnan, Jason Stanley, Omesh Tickoo, Lama Nachman, Rumi Chunara, Madhulika Srikumar, Adrian Weller, and Alice Xiang. 2021. Uncertainty as a Form of Transparency: Measuring, Communicating, and Using Uncertainty. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. Association for Computing Machinery, New York, NY, USA, 401–413.
- [15] Abeba Birhane and Jelle van Dijk. 2020. Robot rights? Let’s talk about human welfare instead. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. Association for Computing Machinery, New York, NY, USA, 207–213.
- [16] Harvard Law Review Editorial Board. 2021. Google LLC v. Oracle America, Inc. <https://harvardlawreview.org/2021/11/google-llc-v-oracle-america-inc/>
- [17] Alexei Botchkarev. 2019. A New Typology Design of Performance Metrics to Measure Errors in Machine Learning Regression Algorithms. *Interdisciplinary Journal of Information, Knowledge, and Management* 14 (2019), 045–076.
- [18] Neil E. Boudette. 2021. Tesla Says Autopilot Makes Its Cars Safer. Crash Victims Say It Kills. <https://www.nytimes.com/2021/07/05/business/tesla-autopilot-lawsuits-safety.html>
- [19] Xavier Bouthillier, César Laurent, and Pascal Vincent. 2019. Unreproducible Research is Reproducible. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 725–734.
- [20] Mark Bovens. 2007. Analysing and assessing accountability: A conceptual framework. *European Law Journal* 13, 4 (2007), 447–468.
- [21] Mark Bovens, Thomas Schillemans, and Robert E Goodin. 2014. Public Accountability. *The Oxford Handbook of Public Accountability* 1, 1 (2014), 1–22.
- [22] Karen L Boyd. 2021. Datasheets for Datasets help ML Engineers Notice and Understand Ethical Issues in Training Data. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–27.
- [23] Myles Brand. 1984. *Intending and Acting: Toward a Naturalized Action Theory*. MIT Press, Cambridge, MA, USA.
- [24] Joanna J Bryson, Mihailis E Diamantis, and Thomas D Grant. 2017. Of, for, and by the people: the legal lacuna of synthetic persons. *Artificial Intelligence and Law* 25, 3 (2017), 273–291.
- [25] Ryan Calo. 2015. Robotics and the Lessons of Cyberlaw. *California Law Review* 103, 3 (2015), 513–563.
- [26] Ryan Calo. 2021. Modeling Through. *Duke Law Journal* 72 (2021), 28 pages. <https://ssrn.com/abstract=3939211> SSRN Preprint.
- [27] Michael Carbin. 2019. Overparameterization: A connection between software 1.0 and software 2.0. In *3rd Summit on Advances in Programming Languages (SNAPL 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 13 pages.
- [28] Jonathan Chang, Sean Gerrish, Chong Wang, Jordan Boyd-Graber, and David Blei. 2009. Reading Tea Leaves: How Humans Interpret Topic Models. In *Advances in Neural Information Processing Systems*, Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta (Eds.), Vol. 22. Curran Associates, Inc., Red Hook, NY, USA.
- [29] Danielle K. Citron and Ryan Calo. 2020. The Automated Administrative State: A Crisis of Legitimacy. , 51 pages. https://scholarship.law.bu.edu/faculty_scholarship/838 Working Paper.
- [30] Danielle Keats Citron and Daniel J. Solove. 2022. Privacy Harms. *Boston University Law Review* 102 (2022), 62 pages.
- [31] Ignacio Cofone and Katherine J. Strandburg. 2019. Strategic Games and Algorithmic Secrecy. *McGill Law Journal* 623 (2019), 41 pages.
- [32] A. Feder Cooper and Ellen Abrams. 2021. Emergent Unfairness in Algorithmic Fairness-Accuracy Trade-Off Research. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. Association for Computing Machinery, New York, NY, USA, 46–54.
- [33] A. Feder Cooper, Karen Levy, and Christopher De Sa. 2021. Accuracy-Efficiency Trade-Offs and Accountability in Distributed ML Systems. In *Equity and Access in Algorithms, Mechanisms, and Optimization*. Association for Computing Machinery, New York, NY, USA, Article 4, 11 pages.
- [34] A. Feder Cooper, Yucheng Lu, Jessica Zosa Forde, and Christopher De Sa. 2021. Hyperparameter Optimization Is Deceiving Us, and How to Stop It. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc., Red Hook, NY, USA, 43 pages.
- [35] Kate Crawford and Jason Schultz. 2014. Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms. *B.C. Law Review* 55 (2014), 93–128, Issue 93.
- [36] Donald Davidson. 1963. Actions, Reasons, and Causes. *The Journal of Philosophy* 60, 23 (1963), 685–700.
- [37] Christopher M De Sa. 2020. Random Reshuffling is Not Always Better. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., Red Hook, NY, USA, 5957–5967.
- [38] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186.
- [39] Nicholas Diakopoulos. 2020. Accountability, Transparency, and Algorithms. *The Oxford Handbook of Ethics of AI* 17, 4 (2020), 197.
- [40] Finale Doshi-Velez and Been Kim. 2018. Considerations for Evaluation and Generalization in Interpretable Machine Learning. In *Explainable and Interpretable Models in Computer Vision and Machine Learning*, Hugo Jair Escalante, Sergio Escalera, Isabelle Guyon, Xavier Baró, Yağmur Güçlütürk, Umut Güçlü, and Marcel van Gerven (Eds.). Springer International Publishing, 3–17.
- [41] Cynthia Dwork. 2006. Differential Privacy. In *Automata, Languages and Programming*, Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener (Eds.). Springer, Berlin, Heidelberg, 1–12.
- [42] Madeleine Clare Elish. 2019. Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction. *Engaging Science, Technology, and Society* (2019), 29 pages.
- [43] European Parliament Committee on Legal Affairs. 2017. Report with Recommendations to the Commission on Civil Law Rules on Robotics.
- [44] Federal Trade Commission. 2014. A Call For Transparency and Accountability: A Report of the Federal Trade Commission. , 110 pages.
- [45] Joel Feinberg. 1968. Collective Responsibility, In Sixty-Fifth Annual Meeting of the American Philosophical Association Eastern Division. *The Journal of Philosophy* 65, 21, 674–688.
- [46] Joel Feinberg. 1970. *Doing & Deserving: Essays in the Theory of Responsibility*. Princeton University Press, Princeton, NJ, USA.
- [47] Joel Feinberg. 1985. Sua Culpa. In *Ethical Issues in the Use of Computers*. Princeton University Press, Princeton, NJ, USA, 102–120.
- [48] Michael Feldman, Sorelle A. Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and Removing Disparate Impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (Sydney, NSW, Australia) (KDD ’15)*. Association for Computing Machinery, New York, NY, USA, 259–268.
- [49] Benjamin Fish, Jeremy Kun, and Ádám Dániel Lelkes. 2016. A Confidence-Based Approach for Balancing Fairness and Accuracy. Preprint.
- [50] Jessica Zosa Forde, A. Feder Cooper, Kweku Kwegyir-Aggrey, Chris De Sa, and Michael L. Littman. 2021. Model Selection’s Disparate Impact in Real-World Deep Learning Applications. <https://arxiv.org/abs/2104.00606>
- [51] Harry G Frankfurt. 2019. *Alternate Possibilities and Moral Responsibility/Alternatieve Möglichkeiten und moralische Verantwortung (English/Deutsch.) Reclam Great Papers Philosophie*. Reclam Verlag, Ditzingen, Germany.
- [52] Alex A. Freitas. 2014. Comprehensible Classification Models: A Position Paper. *SIGKDD Explor. Newsl.* 15, 1 (March 2014), 1–10.
- [53] Batya Friedman and Helen Nissenbaum. 1996. Bias in Computer Systems. *ACM Trans. Inf. Syst.* 14, 3 (July 1996), 330–347.
- [54] Jeanne C Fromer. 2019. Machines as the new Oompa-Loompas: trade secrecy, the cloud, machine learning, and automation. *NYU L. Rev.* 94 (2019), 706.
- [55] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé Iii, and Kate Crawford. 2021. Datasheets for datasets. *Commun. ACM* 64, 12 (2021), 86–92.
- [56] Alvin I Goldman. 2015. *Theory of Human Action*. Princeton University Press, Princeton, NJ, USA.
- [57] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and Harnessing Adversarial Examples. In *ICLR (Poster)*, 11 pages.
- [58] Joseph Y. Halpern and Judea Pearl. 2005. Causes and Explanations: A Structural-Model Approach. Part I: Causes. *British Journal for the Philosophy of Science* 56, 4 (2005), 843–887.
- [59] Ronan Hamon, Henrik Junklewitz, Gianclaudio Malgieri, Paul De Hert, Laurent Beslay, and Ignacio Sanchez. 2021. Impossible Explanations? Beyond explainable AI in the GDPR from a COVID-19 use case scenario. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM, New York, NY, USA, 549–559.
- [60] Margot J. Hanley, Apoorv Khandelwal, Hadar Averbuch-Elor, Noah Snively, and Helen Nissenbaum. 2020. An Ethical Highlighter for People-Centric Dataset

- Creation. ArXiv Preprint.
- [61] Moritz Hardt, Eric Price, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (Eds.), Vol. 29. Curran Associates, Inc., Red Hook, NY, USA.
- [62] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference and Prediction* (2 ed.). Springer, USA.
- [63] Deborah Hellman. 2021. Big Data and Compounding Injustice. , 18 pages. Forthcoming, SSRN preprint.
- [64] Kenneth Einar Himma. 2009. Artificial agency, consciousness, and the criteria for moral agency: What properties must an artificial agent have to be a moral agent? *Ethics and Information Technology* 11, 1 (2009), 19–29.
- [65] Wei Hu, Zhiyuan Li, and Dingli Yu. 2020. Understanding Generalization of Deep Neural Networks Trained with Noisy Labels. In *ICLR*. 13 pages.
- [66] Ben Hutchinson, Andrew Smart, Alex Hanna, Emily Denton, Christina Greer, Oddur Kjartansson, Parker Barnes, and Margaret Mitchell. 2021. Towards accountability for machine learning datasets: Practices from software engineering and infrastructure. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM, New York, NY, USA, 560–575.
- [67] Matthias Jarke. 1998. Requirements Tracing. *Commun. ACM* 41, 12 (Dec. 1998), 32–36.
- [68] Muhammad Atif Javed and Uwe Zdun. 2014. A systematic literature review of traceability approaches between software architecture and source code. In *Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering - EASE '14*. ACM Press, London, England, United Kingdom, 1–10.
- [69] Severin Kaciánka and Alexander Pretschner. 2021. Designing Accountable Systems. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) (FAccT '21). Association for Computing Machinery, New York, NY, USA, 424–437.
- [70] Nathan Kallus and Angela Zhou. 2018. Residual Unfairness in Fair Machine Learning from Prejudiced Data. arXiv:1806.02887
- [71] Margot E. Kaminski. 2019. Binary Governance: Lessons from the GDPR's Approach to Algorithmic Accountability. *S. Cal. L. Rev.* 92 (2019), 89 pages. Issue 1529.
- [72] Sunny Seon Kang. 2020. Algorithmic accountability in public administration: The GDPR paradox. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. ACM, New York, NY, USA, 32–32.
- [73] Been Kim and Finale Doshi-Velez. 2021. Machine Learning Techniques for Accountability. *AI Magazine* 42, 1 (April 2021), 47–52.
- [74] Alexandra Kleeman. 2016. Cooking with Chef Watson, I.B.M.'s Artificial-Intelligence App. *The New Yorker* (20 Nov. 2016).
- [75] Jon Kleinberg, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. 2018. Human Decisions and Machine Predictions. *The Quarterly Journal of Economics* 133, 1 (2018), 237–293.
- [76] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, Tony Lee, Etienne David, Ian Stavness, Wei Guo, Berton Earnshaw, Imran Haque, Sara M Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. 2021. WILDS: A Benchmark of in-the-Wild Distribution Shifts. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 5637–5664.
- [77] 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*, Vol. 7. ACM, New York, NY, USA, 7. Tutorial.
- [78] Simon Kornblith, Jonathon Shlens, and Quoc V. Le. 2019. Do Better ImageNet Models Transfer Better?. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 24 pages.
- [79] Jeff Kosseff. 2022. A User's Guide to Section 230, and a Legislator's Guide to Amending It (or Not). *Berkeley Technology Law Journal* 37 (2022), 40 pages. Issue 2.
- [80] Sarah Kreps, R. Miles McCain, and Miles Brundage. 2020. All the News That's Fit to Fabricate: AI-Generated Text as a Tool of Media Misinformation. *Journal of Experimental Political Science* (2020), 1–14.
- [81] Joshua A. Kroll. 2021. Outlining Traceability: A Principle for Operationalizing Accountability in Computing Systems. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) (FAccT '21). Association for Computing Machinery, New York, NY, USA, 758–771.
- [82] Joshua A. Kroll, Joanna Huey, Solon Barocas, Edward W. Felten, Joel R. Reidenberg, David G. Robinson, and Harlan Yu. 2017. Accountable Algorithms. *University of Pennsylvania Law Review* 165 (2017), 633–705. Issue 633.
- [83] Sarah Lambdan. 2019. When Westlaw Fuels ICE Surveillance: Legal Ethics in the Era of Big Data Policing. *N.Y.U. Review of Law and Social Change* 43 (2019), 255–293. Issue 2.
- [84] David Lehr and Paul Ohm. 2017. Playing with the Data: What Legal Scholars Should Learn About Machine Learning. *U.C. Davis Law Review* 51 (2017), 653–717.
- [85] Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, and Stephen Wolff. 2009. A Brief History of the Internet. *SIGCOMM Comput. Commun. Rev.* 39, 5 (oct 2009), 22–31.
- [86] Mark A. Lemley and Bryan Casey. 2019. Remedies for Robots. *The University of Chicago Law Review* 86, 5 (2019), 1311–1396.
- [87] Karen EC Levy. 2014. Intimate Surveillance. *Idaho L. Rev.* 51 (2014), 679.
- [88] Karen EC Levy and David Merritt Johns. 2016. When open data is a Trojan Horse: The weaponization of transparency in science and governance. *Big Data & Society* 3, 1 (2016), 6 pages.
- [89] Charles Lovering, Rohan Jha, Tal Linzen, and Ellie Pavlick. 2021. Predicting Inductive Biases of Pre-Trained Models. In *International Conference on Learning Representations*.
- [90] Donald MacKenzie. 2001. *Mechanizing Proof: Computing, Risk, and Trust*. MIT Press, Cambridge, MA, USA.
- [91] Angelina McMillan-Major, Salomey Osei, Juan Diego Rodriguez, Pawan Sasanka Ammanamanchi, Sebastian Gehrmann, and Yacine Jernite. 2021. Reusable Templates and Guides For Documenting Datasets and Models for Natural Language Processing and Generation: A Case Study of the HuggingFace and GEM Data and Model Cards. ArXiv preprint.
- [92] Alexander Meinke and Matthias Hein. 2019. Towards neural networks that provably know when they don't know. ArXiv preprint.
- [93] Jacob Metcalf, Emanuel Moss, Elizabeth Anne Watkins, Ranjit Singh, and Madeleine Clare Elish. 2021. Algorithmic impact assessments and accountability: The co-construction of impacts. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM, New York, NY, USA, 735–746.
- [94] 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*. ACM, New York, NY, USA, 220–229.
- [95] Michael Moore. 2019. Causation in the Law. In *The Stanford Encyclopedia of Philosophy* (Winter 2019 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University, USA.
- [96] Emanuel Moss, Elizabeth Anne Watkins, Ranjit Singh, Madeleine Clare Elish, and Jacob Metcalf. 2021. Assembling Accountability: Algorithmic Impact Assessment for the Public Interest. SSRN preprint.
- [97] D.K. Mulligan and K.A. Bamberger. 2018. Saving governance-by-design. *California Law Review* 106 (June 2018), 697–784.
- [98] Deirdre K. Mulligan and Kenneth A. Bamberger. 2019. Procurement As Policy: Administrative Process for Machine Learning. *Berkeley Technology Law Journal* 34 (4 Oct. 2019), 771–858.
- [99] Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. StereoSet: Measuring stereotypical bias in pretrained language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, Online, 5356–5371.
- [100] National Highway Traffic Safety Administration. 2016. *Federal Automated Vehicles Policy: Accelerating the Next Revolution In Roadway Safety*. Technical Report. U.S. Department of Transportation.
- [101] Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. 2017. Exploring Generalization in Deep Learning. In *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc., Red Hook, NY, USA.
- [102] Helen Nissenbaum. 1996. Accountability in a computerized society. *Science and engineering ethics* 2, 1 (1996), 25–42.
- [103] Ngozi Okidebe. 2022. Discredited Data. Forthcoming, Cornell Law Review, Vol. 107, shared privately with the authors.
- [104] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, D. Sculley, Sebastian Nowozin, Joshua V. Dillon, Balaji Lakshminarayanan, and Jasper Snoek. 2019. Can You Trust Your Model's Uncertainty? Evaluating Predictive Uncertainty under Dataset Shift. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. Curran Associates Inc., Red Hook, NY, USA, Article 1254, 12 pages.
- [105] Nicolas Papernot, Patrick D. McDaniel, Somesh Jha, Matt Fredrikson, Z. Berkay Celik, and Ananthram Swami. 2016. The Limitations of Deep Learning in Adversarial Settings. In *1st IEEE European Symposium on Security & Privacy*. IEEE, 16 pages.
- [106] Samir Passi and Solon Barocas. 2019. Problem Formulation and Fairness. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (Atlanta, GA, USA) (FAccT '19). Association for Computing Machinery, New York, NY, USA, 39–48.

- [107] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., Red Hook, NY, USA, 8024–8035.
- [108] Alisha Pradhan, Leah Findlater, and Amanda Lazar. 2019. “Phantom Friend” or “Just a Box with Information” Personification and Ontological Categorization of Smart Speaker-based Voice Assistants by Older Adults. In *Proceedings of the ACM on Human-Computer Interaction*, Vol. 3. ACM, New York, NY, USA, 1–21.
- [109] Edward Raff. 2019. A Step toward Quantifying Independently Reproducible Machine Learning Research. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. Curran Associates Inc., Red Hook, NY, USA, Article 492, 11 pages.
- [110] Inioluwa Deborah Raji, Andrew Smart, Rebecca N White, Margaret Mitchell, Timmit 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*. ACM, New York, NY, USA, 33–44.
- [111] Andrew L. Russell. 2014. *Open Standards and the Digital Age: History, Ideology, and Networks*. Cambridge University Press, Cambridge, UK.
- [112] Jathan Sadowski, Salomé Viljoen, and Meredith Whittaker. 2021. Everyone should decide how their digital data are used – Not just tech companies.
- [113] Thomas Scanlon. 2000. *What We Owe to Each Other*. Belknap Press, Cambridge, MA, USA.
- [114] Markus Schlosser. 2019. Agency. In *The Stanford Encyclopedia of Philosophy* (Winter 2019 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University, USA.
- [115] Andrew D Selbst. 2021. An Institutional View Of Algorithmic Impact Assessments. *Harvard Journal of Law & Technology* 35 (2021), 75 pages. Issue 117.
- [116] 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* (Atlanta, GA, USA) (FAT* '19). Association for Computing Machinery, New York, NY, USA, 59–68.
- [117] Hetan Shah. 2018. Algorithmic Accountability. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376, 2128 (2018), 6 pages.
- [118] Catherine M Sharkey. 2016. Can Data Breach Claims Survive the Economic Loss Rule. *DePaul Law Review* 66 (2016), 339.
- [119] Hong Shen, Wesley H Deng, Aditi Chattopadhyay, Zhiwei Steven Wu, Xu Wang, and Haiyi Zhu. 2021. Value Cards: An Educational Toolkit for Teaching Social Impacts of Machine Learning through Deliberation. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM, New York, NY, USA, 850–861.
- [120] David Shoemaker. 2011. Attributability, answerability, and accountability: Toward a wider theory of moral responsibility. *Ethics* 121, 3 (2011), 602–632.
- [121] Prabhu Teja Sivaprasad, Florian Mai, Thijs Vogels, Martin Jaggi, and François Fleuret. 2020. Optimizer Benchmarking Needs to Account for Hyperparameter Tuning. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 9036–9045.
- [122] Mona Sloane, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 2020. Participation is not a Design Fix for Machine Learning. ArXiv preprint.
- [123] Mona Sloane, Emanuel Moss, and Rumman Chowdhury. 2021. A Silicon Valley Love Triangle: Hiring Algorithms, Pseudo-Science, and the Quest for Auditability. ArXiv preprint.
- [124] Brian Cantwell Smith. 1985. The Limits of Correctness. *SIGCAS Comput. Soc.* 14,15, 1,2,3,4 (Jan. 1985), 18–26.
- [125] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research* 15, 56 (2014), 1929–1958.
- [126] Luke Stark and Jevan Hutson. 2021. Physiognomic Artificial Intelligence. SSRN preprint.
- [127] Tony Sun, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang. 2019. Mitigating Gender Bias in Natural Language Processing: Literature Review. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Florence, Italy, 1630–1640.
- [128] Harry Surden and Mary-Anne Williams. 2016. Technological Opacity, Predictability, and Self-Driving Cars. *Cardozo Law Review* 38 (2016), 121–181.
- [129] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. 2014. Intriguing properties of neural networks. In *ICLR (Poster)*, 10 pages.
- [130] Matthew Talbert. 2019. Moral Responsibility. In *The Stanford Encyclopedia of Philosophy* (Winter 2019 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University, USA.
- [131] Piotr Tereszkievicz. 2018. Digital platforms: regulation and liability in the EU law. *European Review of Private Law* 26, 6 (2018), 18 pages.
- [132] Sherry Turkle. 2005. *The second self: Computers and the human spirit*. MIT Press, Cambridge, MA, USA.
- [133] U.S. Government Accountability Office. 2021. Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities. , 112 pages. <https://www.gao.gov/assets/gao-21-519sp.pdf>
- [134] Kees-Jan van Dorp. 2002. Tracking and tracing: a structure for development and contemporary practices. *Logistics Information Management* 15, 1 (March 2002), 24–33.
- [135] Briana Vecchione, Karen Levy, and Solon Barocas. 2021. Algorithmic Auditing and Social Justice: Lessons from the History of Audit Studies. In *Equity and Access in Algorithms, Mechanisms, and Optimization*. ACM, New York, NY, USA, 1–9.
- [136] Carissa Véliz. 2021. Moral zombies: why algorithms are not moral agents. *AI & SOCIETY* 36 (2021), 11 pages.
- [137] Salomé Viljoen. 2021. A Relational Theory of Data Governance. *Yale Law Journal* 131, 2 (2021), 82 pages.
- [138] Sandra Wachter and Brent Mittelstadt. 2019. A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI. *Columbia Business Law Review* 2 (2019), 130 pages.
- [139] Sandra Wachter, Brent Mittelstadt, and Luciano Floridi. 2017. Transparent, explainable, and accountable AI for robotics. *Science (Robotics)* 2, 6 (2017).
- [140] Ari Ezra Waldman. 2019. Power, Process, and Automated Decision-Making. *Fordham Law Review* 88 (2019), 21 pages.
- [141] Yilun Wang and Michal Kosinski. 2018. Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *Journal of Personality and Social Psychology* 114 (02 2018), 246–257.
- [142] Gary Watson. 1996. Two Faces of Responsibility. *Philosophical Topics* 24, 2 (1996), 227–248.
- [143] Max Weber. 1978. *Max Weber: selections in translation*. Cambridge University Press, Cambridge, UK.
- [144] Jan Whittington, Ryan Calo, Mike Simon, Jesse Woo, Meg Young, and Peter Schmiedeskamp. 2015. Push, pull, and spill: A transdisciplinary case study in municipal open government. *Berkeley Technology Law Journal* 30, 3 (2015), 1899–1966.
- [145] 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*. ACM, New York, NY, USA, 1–18.
- [146] Xiaolin Wu and Xi Zhang. 2016. Automated Inference on Criminality using Face Images. ArXiv preprint.
- [147] Yichen Yang and Martin Rinard. 2019. Correctness Verification of Neural Networks. NeurIPS 2019 Workshop on Machine Learning with Guarantees.
- [148] Meg Young, Luke Rodriguez, Emily Keller, Feiyang Sun, Boyang Sa, Jan Whittington, and Bill Howe. 2019. Beyond open vs. closed: Balancing individual privacy and public accountability in data sharing. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM, New York, NY, USA, 191–200.
- [149] Du Zhang and Jeffrey JP Tsai. 2003. Machine learning and software engineering. *Software Quality Journal* 11, 2 (2003), 87–119.
- [150] Ruqi Zhang, A. Feder Cooper, and Christopher M De Sa. 2020. Asymptotically Optimal Exact Minibatch Metropolis-Hastings. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., Red Hook, NY, USA, 19500–19510.