Representativeness in Statistics, Politics, and Machine Learning

Kyla Chasalow
Cornell University
kec89@cornell.edu

Karen Levy
Cornell University
karen.levy@cornell.edu

ABSTRACT

Representativeness is a foundational yet slippery concept. Though familiar at first blush, it lacks a single precise meaning. Instead, meanings range from typical or characteristic, to a proportionate match between sample and population, to a more general sense of accuracy, generalizability, coverage, or inclusiveness. Moreover, the concept has long been contested. In statistics, debates about the merits and methods of selecting a representative sample date back to the late 19th century; in politics, debates about the value of likeness as a logic of political representation are older still. Today, as the concept crops up in the study of fairness and accountability in machine learning, we need to carefully consider the term's meanings in order to communicate clearly and account for their normative implications. In this paper, we ask what representativeness means, how it is mobilized socially, and what values and ideals it communicates or confronts. We trace the concept's history in statistics and discuss normative tensions concerning its relationship to likeness, exclusion, authority, and aspiration. We draw on these analyses to think through how representativeness is used in FAccT debates, with emphasis on data, shift, participation, and power.

CCS CONCEPTS

- **Applied computing** → **Law, social and behavioral sciences;**
- **Mathematics of computing** → **Probability and statistics.**

KEYWORDS

representativeness, sampling, fairness, bias, participation, inclusion

ACM Reference Format:

Kyla Chasalow and Karen Levy. 2021. Representativeness in Statistics, Politics, and Machine Learning. In *Conference on Fairness, Accountability, and Transparency (FAccT '21)*, March 3–10, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3442188.3445872>

1 INTRODUCTION

In the study of fairness, accountability, and transparency in machine learning, there have been various attempts to catalogue and unpack different understandings of the terms “fairness” and “bias” [113, 116, 118, 151]. Less attention has been afforded to “representative,” a foundational yet ambiguous word peppered through discussions of data and statistics. Newspapers allow “nationally representative” surveys to speak for the state of public opinion

and society. A machine learning paper speaks of selecting “representative examples” to further interpretability [133]. Another warns of costs to “representative and consultative” participation when ML technologies scale across contexts [144]. Meanwhile, a legal scholar warns against relying on statistical models developed on “unrepresentative” reference data to make sentencing decisions [68]. Another describes systematic “underrepresentation” in large datasets used to allocate goods and services as a rising form of civic and political exclusion [96].

Linguistic precision is key to meaningful debate, yet some words at the heart of contemporary discussions of data, algorithms, and their consequences resist a single precise definition. Words like “fair,” “biased,” “interpretable,” and “representative” are valuable yet challenging for their multiple meanings and ability to capture intuitive concerns about technical topics. Lipton and Steinhardt warn against the overuse in machine learning of “suitcase words” (a term coined by Minsky) which “pack together a variety of meanings” and have colloquial appeal [99, 111]. While suitcase words can be useful for overarching ideas, they require careful unpacking to make it clear when we are talking past one another or eliding distinctions that have significant conceptual and normative consequences. In this paper, we unpack representativeness.

Fundamentally, representativeness concerns the ability of one thing to stand for another—a sample for a population, an instance for a category. We could think of all of statistics along these lines: as a web of philosophies and practices for describing and making inferences about the world through inevitably limited data. However, unlike many terms of statistical inference—multicollinearity, maximum likelihood, posterior means—“representative” and its complement, “unrepresentative,” appear intuitive and familiar. They engage people in thinking about the trustworthiness and quality of data and the inferences based on them.

In fact, representativeness has long been a contested concept in both statistics and politics. In statistics, debates about the merits and methods of selecting a representative sample date back at least to the late 19th century. As Kruskal and Mosteller have traced, the term has been persistently slippery [89–93]. Its uses within statistics have not been siloed from its more colloquial use to mean typicality, a miniature, the “absence of selective forces,” “general acclaim” for data, or some notion of coverage [90]. In politics, the adjective and noun “representative” is older still, developing in the 17th century the sense of “standing for others” that has become central to representative government [157]. There is an evident connection to representation, another tricky word notably explored by Pitkin in *The Concept of Representation* (1967) [126]. Pitkin locates “representativeness” in the logic of descriptive likeness. This notion of political representation has some affinity with the statistical sample, a connection made explicit in random selection of jury pools, some forms of deliberative “minipublics,” and occasional calls for random sampling of government representatives [36, 58, 126]. It

This work is licensed under a Creative Commons Attribution International 4.0 License.

FAccT '21, March 3–10, 2021, Virtual Event, Canada
ACM ISBN 978-1-4503-8309-7/21/03.
<https://doi.org/10.1145/3442188.3445872>

also, however, raises a persistent debate about the value of likeness relative to other logics of political legitimacy [104, 126].

Since the 1970s, psychologists have also studied representativeness as part of how we think and judge. Tenenbaum and Griffiths call it a “central explanatory construct” used to explain “phenomena of categorization, comparison, and inference,” including “errors in probabilistic reasoning” [152]. The latter refers to the representativeness heuristic, introduced by Tversky and Kahneman to describe our tendency to make judgements based on *similarity* in ways that violate the rules of probability [154, 155]. Other work has focused on understanding our tendency to view some objects as more typical or as better examples of a category than others [7, 108]. But while that work examines representativeness as a cognitive process, representativeness is also an analytic criterion and demand in statistics and politics. We focus on that criterion here.

In this paper, we explore uses of “representative” in the past and present to make sense of its meanings and their relevance for fairness, transparency, and accountability in machine learning. In the words of C.S. Lewis, we look to meet the word “alive” [97, p. 2]. We ask not only what representativeness means but how it is mobilized socially and what values and ideals it communicates or confronts. Our goal is not to pick a best meaning or propose a technical formulation nor to pose representativeness as a panacea. Still, we must be careful when we speak of representativeness, as different meanings have distinct practical and normative implications.

We proceed as follows. In section 2, we trace the history of representative sampling, from early debates about partial data to the emergence of probability sampling. In section 3, we broaden our scope to other scientific contexts, where the word’s historical link to probability sampling provokes debate about different logics of generalizability. In section 4, we argue that representativeness raises normative tensions concerning likeness, exclusion, authority, and aspiration. We use historical examples to illustrate these concepts before pivoting to the present. In section 5, we consider representativeness in FAccT as it connects to data, shift, participation, and power.¹

We see this effort as a complement to other work in FAccT that draws on fields like STS to interrogate taken-for-granted terms and categories [70, 76]. The notion of representativeness presupposes classification and equivalence—that is, the ability for some units to be treated as interchangeable with others, at least in a given context for some salient characteristic [46]. Those equivalences are essential to our ability to learn about the world, but they are also approximate, value-laden, and at times, contested [25, 44]. Only by considering the ways we translate the world into data in the first place can we fully examine the realities data represent and potentially, perpetuate.

¹We acknowledge our own representativeness concerns. We contribute a sense of the variety of meanings and values associated with representativeness, yet do not claim to analyze an exhaustive or representative set of them. We focus more on social sciences than natural sciences and on Europe and the United States more than other parts of the world. Beyond these scope restrictions, issues related to our selection of sources are harder to pinpoint. We take some refuge in being illustrative, yet this does not fully resolve the issue, for what we choose to include—or not—will always shape the reader’s sense of the subject.

2 THE STATISTICAL ROOTS OF REPRESENTATIVENESS

The roots of representative sampling lie in inductive inference. As Stephan wrote in a 1948 history, “All empirical knowledge is, in a fundamental sense, derived from incomplete or imperfect observation and is, therefore, a sampling experience” [149]. But while the act of drawing conclusions from partial information and assessing their validity is as old as the human condition, formal probability-based sampling methods were novel at the turn of the 20th century.

That is not to say that pre-20th century quantitative work was devoid of partial data and inference. Lacking census data, 17th and 18th century political arithmeticians extrapolated from local records and their own surveys for insurance and policy purposes [127, 149], using, for example, data on local birth rates to estimate total population size [47]. These forays into inference from partial data were controversial and remained so into the 19th century, when the growth of statistical bureaucracies and regular census-taking in many nations shifted the emphasis of the emerging field of statistics from partial surveys to complete enumeration [47, 65].

Nineteenth century statistics also featured a tension between notions of statistics as a descriptive science and as a science for discovering laws and causes. While adherents of the former emphasized “one-to-one correspondence between observations and numbers,” with “numbers... interpreted as summaries of immediately observable facts,” those working in the latter framework began to calculate aggregates and make inferences and predictions [138, p. 47–51]. Though not discussed as such at the time, we might view these tensions as early debates about the possibility for representativeness—for a move beyond that one-to-one correspondence. They also provide context for why, when representative sampling was first formally proposed to the International Statistical Institute (ISI) by statistician Anders Nicolai Kiaer in 1895, some viewed the prospect as “dangerous,” with one critic opining that “one cannot replace by calculation the real observation of facts” [93].

2.1 Representativeness as Design Problem and Sociopolitical Problem

By the turn of the 20th century, representative sampling began to be conceptualized as a more formal design problem—an issue to be approached via “systematically designed plans” for data collection [84, p. 37]. Kiaer presented his experiments with sampling as evidence that it was possible to obtain an “approximate miniature” of the population, capturing not just “average conditions” but “information on variation and extreme values” in a rigorous way [84, p. 53]. Historians and statisticians have distinguished the subsequent development of sampling from earlier uses of partial data by its formal procedure and its application of probability theory [140, 149]. Yet part of what was new in this time was also the speed, type, and detail of information that sampling began to provide. “By shrinking drastically the delay between request and delivery of information,” Prévost and Beaud write, sampling “allowed for the production of data that would have had no meaning in the old time frame” [130, p. 154]. Together with the growth of statistical bureaucracies and new mechanical tabulation techniques for census-type data, these data expanded the “knowledge-gathering capacities of modern states.”

However, the representative method itself also emerged from changing economic and political demands on statistical “descriptions of the world” and related changes in thinking about parts and wholes [47, p. 211]. While earlier monography work intended to illustrate “holistic” truths by selecting the average type, sampling emerged from a need for information on *variation* among individuals [47, 109]. As local charity solutions to poverty became inadequate in the context of national economic crises, surveys needed to provide not holistic knowledge of the typical family but nationally applicable knowledge of the condition of individuals so that governments could allocate and distribute aid. In the American context, the shock of the Great Depression “ruptured” established links between “statistics and political action” and created a need for a “new numerical reading of America” that could support a more interventionist government [48, p. 4, p. 81]. Businesses and election campaigns with expanding reach also demanded national-level information at low cost and high speed. Thus, as Desrosières argues, debates about how to obtain a representative sample of a population could only be formulated once “the problem itself... the constraint of representativeness” became salient [47, p. 210–1]. Representativeness was not only a design problem; it was also a sociopolitical problem translated into new ways of doing statistics.

2.2 Achieving Representativeness: Mechanical and Purposive Selection

By the 1920s, initial debate among statisticians about the possibility of sampling had largely given way to the question of *how* to sample [18]. Given support for both random and purposive sampling methods, a 1926 ISI report endorsed both [79]. Then, in 1934, a landmark paper by Jerzy Neyman helped establish random over purposive methods, formalized stratified sampling, and contributed to the theoretical foundation of probability sampling [18, 93]. The full story of these ideas and their adoption is more layered, involving interactions among theory, practical experimentation, political demands, and bureaucratic processes [48]; for our purposes, we highlight a few main theoretical tensions about how to achieve—and what it means to be—a representative sample.

“It will be obvious,” Kiaer wrote in 1897, “that the representative method can be applied in several ways” [84, p. 43]. He described two samples, which, though not labeled “mechanical” or “purposive,” foreshadowed these logics. The first sample is premised on arbitrary filtering, including by the first letter of surnames. He reasoned that such selection should occur in a “haphazard or random way” so as “to avoid, in the most stringent manner, any procedure that could give preference to persons in certain occupations or belonging to particular social strata” [84, p. 39]. This mention of randomness and preference suggests a notion of mechanical objectivity [43, 128], in the sense that by following rules in a “stringent manner,” representativeness would result from the absence of selective discretion.

The same appeals underlie random sampling, a bedrock of probability theory. While Kiaer does not mention probability, statisticians would soon begin to emphasize its role and the value of estimating “probable limits of error” [26, 81]. The “first essential” of sampling, wrote Bowley in 1910, is “that every member of the group considered should have nearly the same chance of being included in the sample” [27, p. 56]. It was standard early on to emphasize this

“equal chance” as a defining feature of mechanical approaches [18], perhaps because it suggests a form of procedural fairness (some even spoke of “fair samples” [29, 93, 107, 148]) and neutrality. After all, random samples are only guaranteed to be accurate and non-preferential on average. As Stephan reflected in 1939, “It provides but little aid and comfort to the victim of a poor draw in random sampling to assure him that ‘in the long run’ the random method of selection will give him errors in one direction just as often as it will the others” [147]. The issue of how to check a *particular* sample for representativeness remained a point of debate [62, 79, 93, 107]. But locating representativeness in a mechanical *procedure* did provide a way to move past the specter of chance skew and uncheckable variables. In the words of statistician Margaret Hogg, a firm proponent, “only if a random sample has been secured can one argue from representativeness which has been tested, to representativeness where no test is possible” [75].

The “victim of a poor draw” in random sampling might favor another approach: using available data to ensure a proportionate match on known relevant variables. In Kiaer’s second sample, surveyors in rural areas used census data to allocate counts per county and then selected districts within counties to “represent the main industry-groups within the country as well as its various geographic conditions” [84, p. 41]. On the ground, enumerators were trusted “to select respondents in a representative manner ... [and] check that not only typical middle-class houses were visited but also those occupied by the more well-to-do and the poorer classes” [84, p. 42]. Later, surveyors collected additional data to correct non-matching proportions on variables such as occupation—a safeguard against enumerator bias. It seems to be the match that made the sample “representative,” not the exact selection rule or persons selected.

What about variables for which no matching is possible? Kiaer provides only a general assurance that samples will yield “the same degree of accuracy when it comes to new fields.” He follows with an analogy: “The relationship between these two sets of returns can be illustrated by the relationship between two barrels, one small and the other very large... If a number of samples taken from each, crosswise and lengthwise in every conceivable direction, show the composition of the content of both barrels to be practically the same, I think it is justified, on the basis of further, more extensive analysis that we assume can only be carried out on the smaller barrel, to draw valid inferences about the unknown content of the larger barrel” [84, p. 51]. This barrel analogy suggests what we call the *universal representativeness ideal*—a proportionate match for every variable at once. Its ambition is to scale down without loss of information. Its allure lies not only in a correct answer to an immediate question. Better yet, it promises the possibility of data reuse via a kind of representativeness that holds from every angle.²

Purposive sampling formalized the idea that a sample matching on some variables would match on others; the added ingredient was correlation. Formally, it would come to involve matching to average values of known variables (“controls”) thought to be correlated to unknown variables of interest [79, 80]. What especially set it apart

²Analytically, this ideal is impossible. Not only can we not check every possible variable, but universal representativeness would also have to apply to joint distributions of variables. There is also an infinite regress issue: if the whole contains a miniature, why should not the miniature contain a miniature? Yet at some point, the reduced version would become so small that we would not soon call it representative [90].

from mechanical approaches was its ability to incorporate *existing knowledge*. To repeat a pedagogic example discussed in [119], if we wished to learn the average right arm length of a group of people and we knew the lengths of their left arms, we might select people whose left arms were closest to the average left arm length. The fact that left and right arm lengths are highly correlated would likely lead to a fairly accurate result. In applied problems, correlations are often smaller than left-right arm correlations are, and practitioners who attempted the method encountered problems [40, 93, 119]. It was stratification that would become a foundational method for incorporating existing knowledge into probability sampling.

2.3 Stratification: A Representative Method

Although stratification appeared in some earlier sampling work [28, 109], Neyman helped establish it for a wider audience [18]. Stratification bridged the purposive-random tension, providing a way to increase the precision of estimates by incorporating knowledge of group differences into random sampling [47, 148]. It was also part of the emerging theory of probability sampling, in which unequal chances of selection were allowed and sometimes preferable. Today, a probability sampling design refers to a specification of possible samples and their probabilities of selection. Equal probabilities of inclusion are not required; what matters is that the probabilities of inclusion of each unit are known so that they can be used as weights [100]. The literature that has built on these ideas over the last eighty or so years is expansive, but for our purposes, the interesting question is how these ideas relate to representativeness.

A less discussed feature of Neyman's work is the way that he relocates representativeness to a property of a method. "Obviously," he wrote in 1934, "the problem of the representative method is *par excellence* the problem of statistical estimation," later adding, "If there are difficulties in defining the 'generally representative sample,' I think it is possible to define what should be termed a *representative method of sampling* and a *consistent method of estimation*" [119, p. 572]. By attaching "representative" to "method" and not to "sample," Neyman avoids the uncertainties inherent to any one sample and focuses on building a general theoretical framework for inference that performs well on average and allows estimates of precision. In the process, he broadens the scope of "representativeness" to apply to any legitimate probability-based design—a sample could be disproportionate but result from a representative method.

In practice, people still speak of representative *samples*, but when applied to probability sampling in general, it seems that accurate maps from sample to population are the basic ingredient for representativeness [91, 92]. It then makes sense to speak of a sample being "weighted to be representative" (e.g. [114]), perhaps restoring a sense of proportionality. If, by design or implementation, some subgroups of the target population lack any representatives, if weights are inaccurate, or if other distortions arise that impede inference, then the sample (or method) might be called unrepresentative (e.g. [12, 39]). Admittedly, this is broad, and it remains an open question how sample size, degree of uncertainty, and other methodological elements should factor in here. But it is also flexible. In a way, we have come full circle, back to the basic link between statistical representativeness and scientific inference.

3 REPRESENTATIVENESS IN SCIENCE

Most broadly, representativeness is a generalizability concern. But generalizability is itself a suitcase word, packing in notions of sample-to-population inference, transfer across populations, and external validity from experiments to real-world settings [52, 55, 99, 137]. Within that assortment, some use representativeness primarily to describe sample-to-population inference yielding time- and place-specific descriptions of population frequencies, as in much survey research. Here, the question of whether a sample is representative of a target population may be complemented by a scope of inference question: what are the data at hand actually representative of [92]? In contrast, the question of whether, for example, biological findings in fruit flies apply to humans or of whether the results of experiments with college students apply to other people introduce generalizability logics different from the probability sampling one. It is important to distinguish these logics, but two realities complicate those distinctions.

First, in practice, representativeness is sometimes used more flexibly to describe patterns of exclusion and inclusion across objects of scientific study. In clinical trials, for example, the word is used to describe the over- and under-representation of minorities and other consistent selection biases, which raise both generalizability and equity concerns [31, 56]. Similarly, in genetic research, the fact that people of European ancestry make up a large portion of data used to study disease genetics has consequences both for underrepresented groups and for science [66, 143]. These examples concern (un)representativeness across fields of research and their collections of studies, datasets, and subjects. A related flavor of the concept is used to discuss the "model organism" problem. The term refers to the organisms commonly used in biology, such as *Drosophila melanogaster* (the fruit fly) [21]. It has been extended as a metaphor for when a field relies heavily on certain research objects (Twitter in social media research [153], chess in AI [53]), with consequences for the field's direction and findings. Though often simple and accessible, model organisms and collections of model organisms may, in part because of these traits, exclude certain kinds of organisms relative to the 'population' of possible objects of study. The possibility for such exclusions to shape the outcomes and causal pathways we observe broadens the scope of representativeness, suggesting it can underlie other generalizability logics.

The second complication is that distinctions between types of generalizability are more than semantics—they are sites of methodological contestation. Thus the choice of whether and how to use "representative" is not purely a matter of its technical meaning. For example, a recent debate among epidemiologists about the importance of representativeness in their field turns out to also be a debate about the scope and focus of that field [50, 52, 120, 135, 136]. "Exalted along with motherhood, apple pie and statistical significance," one author writes, "representativeness may be essential for conducting opinion polls, or for public-health applications, but it is not a reasonable aim for a scientific study," adding that the former are "not science in the same way that causal studies about how nature operates are science" [136]. Another paper pushes back, allowing representativeness a somewhat broader scope beyond random sampling but also defending the scientific value of "descriptive epidemiology" [50]. In qualitative research, debate about the use of

“representative” to describe case selection touches on long-running tensions around the infringement of quantitative logics and standards [15, 145]. “Representative” is a bridging word—occurring in many fields and providing a common language for specialist and non-specialist concerns—but it is also used to draw boundaries.

The preceding discussion creates ambiguities concerning how far the word “representative” extends or should extend. On one extreme, it has a fairly specific historical link to probability sampling. On the other, representativeness is flexible, arising in discussions of over-, under-, non-, or mis-representation of descriptive characteristics or causal mechanisms, for particular samples and for collections of many kinds. This variety suggests a need for caution. There is value in the term’s flexibility, but this paper is not a call to pepper our prose with the word—quite the opposite. If anything, researching representativeness has convinced us that we should be sparing in its use, or at least, that “representative” can rarely stand alone.

4 NORMATIVE TENSIONS OF REPRESENTATIVENESS

(Un)representativeness has a fascinating capacity to describe scientific uncertainties, statistical methods, principles of government, and social (in)equities—sometimes together. So far, we have focused on representativeness mainly as a matter of accurate inference and scientific concerns. We now pivot to a more sociopolitical and normative perspective, for representativeness is not only a descriptive fact; it is also a value, ideal, and mark of authority—merited or not. Here, we untangle some of the normative tensions that surround the concept, illustrating them with historical examples, before turning in section 5 to consider representativeness in the FAccT context.

4.1 The Likeness Ideal

Within her taxonomy of representation, Pitkin locates representativeness in descriptive resemblance—in standing “for others ‘by being sufficiently like them’” [126, p. 80]. As we have seen, applied to data, this suggests a proportionate match or perhaps just a known relationship between data and some target of inference. Applied to social and political processes, it suggests a value for likeness between people and their representatives on relevant variables. It is this that gives concerns about the representativeness of these processes some affinity with statistical concepts of representativeness.

Political processes often do not involve everyone. Representativeness promises to reconcile those exclusions with principles of equity [104, p. 130]. For example, the fact that not all people are activists may be more acceptable if activists are deemed “representative” of those who do not participate—thereby preserving all individuals’ right to be heard [156]. This promise of restoring a sense of inclusion reflects two logics—one normative and one epistemic [122]. First, representativeness may be a normative end linked to voice and democratic legitimacy, as well as to the idea that those who share relevant identities or experiences with the people they represent are better advocates for those people—what Phillips termed the “politics of presence” [125]. Even when it does not translate to direct action, likeness may have symbolic value, inspiring trust [110]. Second, representativeness may be an epistemic means to inclusive and generalizable knowledge about people’s views and experiences. These epistemic goals may still be motivated by equity

issues and, since knowledge may be tied to identity and experience, are sometimes served by the “presence” logic [71].

While the appeals described above may be more prominent in political contexts, they can also apply to data, particularly when data are a form of sociopolitical representation. In either case, the value of likeness may be contested. In political representation, it confronts other logics of political legitimacy, including those that emphasize “knowledge, expertise, or judgement” or authorization and accountability more than shared identity [122]. Here, representativeness as likeness may be irrelevant or rejected in favor of other goals. As described in section 3, the value of representative data may also be contested. But as we now illustrate with a case from the history of sampling, even when some form of representativeness is the goal, it can be complicated by other values and logics.

In 1862, concerned about farmers’ vulnerability to speculators spreading false information, the *American Agriculturist* published a call for “readers...in every town [to] counsel together, and select some man who may be relied upon for good judgement, and general ability to estimate with some degree of accuracy in regard to the leading crops” so that together, a few hundred or more reports from across the country could give “an approximation to the average of the whole” [82]. This method was soon adopted by the new U.S. Department of Agriculture, which cultivated long-term relationships with a network of volunteer crop reporters valued for their intellect and public spirit [48]. These volunteers were not intended to be *like* other farmers—indeed, they were selected precisely because they were deemed to be *unlike* other farmers in their exceptional discernment of their and their neighbors’ crop yields. Accurate data were to result from a kind of epistemic coverage.

However, in the 1920s, problems with the method’s accuracy—attributed in part to reporters being misled by their neighbors [48]—led to the “individual-farm” approach, in which farmers reported only on their own farms [16]. This transformed farmers from *spokespeople* to *examples*, making it necessary to consider their representativeness [48]. Here, the same exceptional status that had legitimized agricultural data before became an obstacle. The eventual solution was probability sampling. The cost was a new problem of nonresponse, for “respondents could refuse to answer, but could not refuse to be selected” [48, p. 195]. Data’s role as a form of political representation changed, too. Because random selection precluded sustained relationships and because, in the New Deal context, data were meant more for planning government interventions than for farmers’ use, the move engendered more separation between data collector (government) and data subject (farmer). In a sense, “representativity had replaced participation” [48, p. 209].

The case of the crop reporters illustrates a number of tensions. While crop reporting reflects the democratic appeal of participatory data collection, it is also exclusionary, reflecting a political logic that emphasizes knowledge and skill over likeness. In contrast, designed sampling reflects the appeal of purposeful, structured inclusion, but at a cost to voluntary participation. These tensions between design and volunteerism and between representatives selected to be exceptional or broadly inclusive will not be new to social and political scientists—but they manifest anew, for example, in debates about crowdsourcing today [9, 137, 146]. Finally, the case highlights that the inclusiveness of a representative sample is no guarantee of power over how data are used. Broadly, then, representativeness,

rooted in the idea of one thing standing for another by “being like” it, interacts with other values and logics. It is already difficult to define and achieve representativeness, yet it is also crucial to ask what it achieves and whether it is the right goal in the first place.

4.2 Exclusionary Representativeness

Recall the universal representativeness ideal suggested by Kiaer’s barrel analogy: a perfect miniature on every variable. Jorge Luis Borges identifies the impossibility of this ideal in his short story “The Congress” (1977), in which an idealistic group attempts to devise a Congress to represent all people everywhere. The group quickly encounters intractable obstacles in deciding which variables matter: “Alejandro Glencoe might represent not only cattlemen but also Uruguayans, and also humanity’s great forerunners, and also men with red beards, and also those who are seated in armchairs. Nora Erfjord was Norwegian. Would she represent secretaries, Norwegian womanhood, or—more obviously—all beautiful women? Would a single engineer be enough to represent all engineers—including those of New Zealand?” [22]. The variables multiply, and it turns out that the only perfect “Congress of the World” is the world itself. In practice, we must decide what variables are salient to evaluate for their representativeness. Epistemic and normative goals can lead to different criteria here—the former concerned with whether variables correlate with those of interest, the latter concerned with social or political relevance [122]. It is worth emphasizing that in either case, these goals can be shaped by our ideals—and that, despite the link between representativeness and inclusion, such ideals can have exclusionary consequences.

Consider the case of *Middletown*, a 1929 study of Muncie, Indiana by Helen and Robert Lynd, notable at the time for applying an anthropological lens to everyday American life [77, 101]. The over-five-hundred-page book was mostly qualitative, interspersed with snippets of interviews with residents on “getting a living,” “making a home,” “using leisure,” and more, but throughout, occasional tabulated data helped reinforce its status as an objective empirical study. The book was an unexpected bestseller. Though widely received as an unprecedented window into a representative American community, the Lynds were cautious but inconsistent regarding *Middletown*’s representativeness [77]. They warned that “a typical city, strictly speaking, does not exist,” but did write that one of their selection criteria was “that the city be as representative as possible of contemporary American life” [101, p. 3]. Of course, they also called the study “Middletown,” a name which abstracts it from Muncie and suggests the ordinary and average, ideas susceptible to being transmuted into typical and representative.

As Igo has traced, the Lynds then make a puzzling acknowledgment: Muncie is “unusual” [77]. They say they have selected for a city of moderate size with “a small Negro and foreign-born population,” adding that to avoid dealing at once with racial and cultural change, “it seemed a distinct advantage to deal with a homogeneous, native-born population, even though such a population is unusual in an American industrial city” [101, p. 8]. They see no contradiction. Muncie is both “as representative as possible” and unusual. In fact, there was a growing black population in Muncie, but the Lynds excluded them from their results (see e.g. fig. 1) [77, p. 56-7]. We might label *Middletown* unrepresentative of Muncie or

Statement	Percentage answering “True”		Percentage answering “False”		Percentage answering “Uncertain”		Percentage not answering	
	Boys	Girls	Boys	Girls	Boys	Girls	Boys	Girls
The white race is the best race on earth.....	66	75	19	17	14	6	1	2
The United States is unquestionably the best country in the world.....	77	88	10	6	11	5	2	1
Every good citizen should act according to the following statement: “My country—right or wrong!”.....	47	56	40	29	9	10	4	5
A citizen of the United States should be allowed to say anything he pleases, even to advocate violent revolution, if he does no violent act himself.....	20	16	70	75	7	7	3	2
The recent labor government in England was a misfortune for England.....	16	15	38	20	38	57	8	8
The United States was entirely right and England was entirely wrong in the American Revolution.....	30	33	55	40	13	25	2	2

¹⁸ Students were requested to write “true,” “false,” or “uncertain” after each statement. No answers of Negroes are included in this summary. The greater conservatism of the girls in their answers to some of the questions is noteworthy. See Appendix on Method. See Index for other answers.

Figure 1: Though *Middletown* was praised for its objectivity, it was more than a recording of facts. An especially stark example is the inclusion of a question asking schoolchildren about their endorsement of white supremacy while “excluding the answers of Negroes” [101, p. 200].

America, but we should also ask in what way the Lynds could suggest it was representative. The answer, Igo writes, is that “the Lynds’ representative community was less an empirical than a normative proposition” [77]. Their goals were epistemic but their concept of their object of study was shaped by an ideal of a “purer, simpler, even preindustrial, America,” especially the white, Protestant part of it—an ideal more “wished-for than real” [77, p. 59].

Admittedly, *Middletown* is different from the statistical samples we discussed earlier. It is a single case study raising a “representativeness as typicality” logic and ironically, therefore particularly susceptible to being in some way atypical.³ But Igo finds a similar tendency in early Gallup opinion polls, which she suggests under-sampled African Americans, immigrants, women, and people with lower incomes in part because they did not fit pollsters’ normative concept of what the likely voter *should be* (reality did come back to bite in the 1948 election) [77, p. 138]. These cases illustrate a form of representativeness claim defined relative to ideals and prescriptions rather than empirical data alone. This does not render the empirical data useless, but those data do exist within a framework in which some variables and individuals are deemed more relevant than others to what it means to be representative. These considerations are always present to some degree as we confront the immense “Congress of the World.”

4.3 Rhetorical Authority

A third tension surrounding the concept of being “representative” concerns the authority that the word itself can inspire, merited or

³“Typical” is yet another common but ambiguous word. It is sometimes used synonymously with representativeness but has more sense of being characteristic, ideal, or both—concepts that are exclusionary in that they deemphasize variety [43, 89].

not. The Lynds *were* somewhat cautious about making representativeness claims. The public reception to *Middletown* was less so. Those who did not read it themselves were liable to encounter it via the newspaper, church, or community event, where its findings were received as if they reflected the state of America [1–5]. The *St. Cloud Times* of Minnesota opined that: “The picture of ‘Middle-town’ is representative of millions of our population, groping for a national culture. The social surveys are unbiased and scientific. It is a book, which should be read by all thoughtful Americans” [6]. Such claims seem to be less about specific numbers or descriptions and more about a vague sense that *Middletown* captured “us” (or at least, those of “us” who fit its ideals) [77]. This public reaction reflects the way representativeness can take on a life of its own, capturing the imagination while also deriving authority from its apparent status as an objective scientific property. This poses a challenge for communication. Words can travel beyond their original intent. Mundane and resonant, imprecise yet somewhat technical, “representative” may be especially liable to do so.

Figure 2: The headline of a 1938 *New York Times* article [60].

Moreover, the word’s authority can be wielded strategically. We see this, for example, in the way George Gallup advertised his new opinion polls to the American public in the 1930s and 40s. Gallup did more than present statistical theory. He translated statistical representativeness into political representation, writing of miniature electorates, sampling referenda, and representatives chosen from every group (see fig. 2) and framing the polls as a way to support democracy by restoring the “voice of the common man” [61, p. 13]. At the same time, he emphasized the polls’ scientific legitimacy to distinguish them from the public failure of the *Literary Digest* polls in the 1936 election, a failure he attributed to a misplaced focus on sample size rather than representativeness [61]. “New polls... based on scientific principles” prioritized the “cross section,” enacting the principle that “the most important requirement of any sample is that it be as representative as possible of the entire group or ‘universe’ from which it was taken” [61, p. 57].⁴

The problem with the rhetorical authority of the word “representative” is that it can be unmerited or misleading. It is misleading to present results without addressing their representativeness, but it is also misleading to suggest representativeness without specifying in what way and to what extent (even then, those cautions are sometimes ignored). The word’s rhetorical halo can also mask

⁴Though Gallup spoke of random sampling here, in early years, he used quota sampling, which set target proportions to mirror important divisions in the population of interest and then charged interviewers with finding people to fill the targets.

important dynamics, including, in the case of Gallup polling, the question of who the data in question most benefited. Gallup’s enthusiastic rhetoric of the democratic “sampling referendum” does not account for the reality that polling had close links to market research and served the needs of businesses, media, and political parties—at times more “surveillance” than “voice” [19, 51].

4.4 Time and Aspiration

Data are inherently historical. Barring time travel, it is impossible to collect data on the future. As Nohr and Olsen have written, “Representativeness is time- and place-specific, and will therefore always be a historical concept. Representativeness is gone as we speak, as Heraclitus told us more than 2000 years ago: ‘You cannot step twice into the same river’” [120]. In this spirit, we turn now to questions about representativeness and time.

Bouk describes a late 19th century struggle between African Americans and major American life insurance companies that raises temporal and normative concerns which implicate representativeness [24]. In the 1880s, in the wake of the Civil War and Reconstruction period, thousands of African Americans submitted applications for life insurance. At first, insurance companies were mostly willing to sell. The catch: African Americans were charged the same rate as white applicants but for only two-thirds of the standard benefit. This, the insurers argued, was not a race-based decision but a statistical one. The companies drew on the 1870 census and other data to show that African Americans had higher mortality rates, an unsurprising finding given recent slavery and ongoing discrimination. The insurers also went further, arguing that these differences applied not only to the current time but always. Emphasizing their reliance on the facts, the life insurers draped racial prejudice with a statistical veil. Bouk introduces the term “fatalizing” to describe how insurers used past data to assign African Americans to a statistically backed fate.

Beginning in 1884, activists succeeded in pushing for anti-discrimination laws in a number of Northern states. Their arguments had two bases. First, they opposed fatalizing practices by defending the possibility for mortality rates to change. “Where life insurers initially assumed continuities between the slave-era past and a free future,” Bouk writes, “African Americans and their allies interpreted the Civil War as a moment of rupture, severing the past from the present and future” [24, p. 32-3]. The question was not so much whether the data used by life insurers were accurate at the time they were collected as whether those data could stand for the present and future. Second, however, advocates emphasized equality in insurance as “a matter of right” [24, p. 41]. Even if mortality rates were different at present or if they remained so at some future point, the principle of equality trumped even accurate fatalism.

This case helps us consider the time-dependence of representativeness in three ways. First, consider the link between the past and the present. The case raises an *empirical* concern about change: even data representative in the past may not be representative in the present if the object they describe has changed. But second, what if past data *do* accurately capture the present, but that present still bears the mark of historical discrimination and injustice? Here we confront an *aspirational* tension concerning future possibility and what some term “historical bias”—“a normative concern with

the world as it is” (and how it might be different) that applies even when we sample and measure perfectly [151]. Insurers’ fatalizing claims not only obscured any change that might have already occurred; they also worked against change by reinforcing existing inequities. In prediction especially, where the object is inherently a future event, feedback loops and future possibility raise questions about what it means for past data to, in a sense, ‘stand for’ the future, together with a normative objection to being represented even by data that may, at present, be representative.

Finally, when advocates argued that non-discrimination in insurance was a “matter of right,” they implied that rights and aspirations can be more important than data—representative or not. This echoes our earlier assertion: that representativeness confronts other values, logics, and goals and that the concept’s intuitive link to inclusion is not straightforward. Overall, we emphasize that the time-dependence of representativeness is an empirical and normative matter. The decision to accept past data as representative can be a matter of accuracy—how much has the present situation shifted relative to the past?—but also of aspiration—do the data reflect the future we would like to see?

5 REPRESENTATIVENESS IN FAccT

Recall that formal sampling emerged in the early 20th century not only as a design question posed by statisticians but also as a sociopolitical demand for new kinds of data that people translated into new ways of doing statistics [47]. Similar dynamics emerge today. In machine learning and the FAccT community, we again see social problems translated into new ways of working with data, with the potential both to exacerbate social problems and to understand and address them [8, 86]. Within that space, representativeness is not a new concern, but its relationship to FAccT concepts and values should be carefully considered. Drawing on what we have learned from tracing the concept through statistical history, scientific discourse, and sociopolitical perspectives, we map four areas where the concept arises in FAccT—data, shift, participation, and power.

5.1 Data

At the turn of the 20th century, mechanical tabulation and sampling helped expand government and commercial data collection, creating an influx of data—so much so that in 1905, Mandello worried that the “ever increasing” demand for statistics had yielded “a mass of uncontrollable figures for the uncontrolled use of everybody” [103]. Today, we again live amid an influx, and it has become common to speak of “big data,” characterized by new scale, speed, formats, and computational approaches. While some have hailed a move away from “small data” probability sampling [106], others have defended its importance [42, 115] or emphasized that representativeness-related questions of selection bias, missingness, and generalizability remain crucial for big data [30, 121]. In a similar vein, though the “GIGO principle” (garbage in; garbage out) is well-known, some have argued that we need to do more to evaluate and communicate where datasets used in ML come from [45, 64, 124, 139].

In machine learning, representativeness raises new flavors of ambiguities we encountered earlier. These include the question of whether “representative data” refers to a proportionate match to some target population or to a more flexible sense of being able to

learn “generalizable predictive patterns” [35]. The idea of learning predictive patterns lies at the heart of ML, but to link these patterns to the words “representative” or “generalizable” is a nontrivial step that can obscure distinctions. Two examples illustrate the danger.

First, data summarization algorithms are sometimes described as reducing data to a “representative set” [23, 37, 54]. In an echo of the earlier split between mechanical and purposive (judgement-based) selection, they do not *per se* sample randomly but rather seek to intentionally remove redundancies while preserving relevant information—in effect, to isolate the signal from the noise. Crucially, however, “representative” here concerns a subset of the original dataset and not necessarily that larger dataset relative to some external target—something the word “representative” by itself does not tell us. Second, the goal of distinguishing meaningful patterns from noise also motivates the practice of evaluating models on holdout data to detect overfitting to training data—often described in terms of measuring “generalization error” and improving “generalizability” to other data from the same distribution [23, 99].⁵ Yet we should distinguish that a non-overfit model, while less sensitive to variations in the data, may still capture an unrepresentative (or non-generalizable) overall pattern relative to some *intended* whole or deployment context. There is a risk here that “generalizable” or “representative” appear to promise more than they do—all the more so given that the external target or deployment context is exactly what people often use “representative” to mean, as a kind of synonym for “real-world” testing and applicability (e.g. [11, 124]).

In FAccT, further ambiguities ensue when representativeness mixes with suitcase words such as “fairness” and “bias” and the related phrases “fair representation” and “representation bias.” “Representation bias” can refer to at least two distinct issues: (1) a disproportionate sample relative to a true distribution; or (2) a class imbalance problem, in which smaller groups may face worse model performance even when sampled proportionately [151]. This latter concern is one common notion of unfairness [38], and it suggests that if representativeness means a proportionate match, it might at times yield *unfair* representation. Yet some do use “representative” to discuss equal representation. To that end, Buolamwini and Raji coin the phrase “user-representative sets” to describe datasets that have “equal representation of each distinct subgroup of the user population” (which might involve oversampling smaller groups) [132]. Such data are “representative,” perhaps, in that they ideally contain sufficient data to allow accurate prediction for members of all relevant groups—a kind of coverage. Of course, even this form of representativeness is not the cure for all forms of bias or unfairness, which also depend on modeling decisions and sociopolitical context [63, 78, 132]. Data, while essential and sometimes overlooked, are only one piece of that puzzle.

So far, we have focused on the single dataset, but what was true of clinical trials and model organisms is also true here: representativeness may concern *collections* of research objects. Machine learning and FAccT research itself often rely on various commonly used datasets, which act as benchmarks for comparing models and methods [34, 69, 158]. If benchmarks do not “represent the target population,” this can create “evaluation bias,” rewarding models that

⁵When we split into train and test sets, “we expect that a representative sample will be in each subset” [23], showing how embedded the idea of random selection as representative has become. But “representative” here is relative to the original dataset.

perform well on unrepresentative data [151]. More fundamentally, these datasets form what Denton *et al.* call an “infrastructure” for the field, acting “analogously to model organisms” and becoming “stand-in[s] for more complicated data traces and machine learning tasks” [45]. In the process, they may individually or collectively reproduce selection tendencies as well as classification and measurement decisions across studies and applications (e.g. [142]).

5.2 Shift

A basic assumption in supervised machine learning is that data distributions match across training, testing, and deployment. As Hand writes, “Intrinsic to the classical supervised classification paradigm is the assumption... that future points to be classified are drawn from the same distributions as the design set,” adding pessimistically that “the assumption that the design distribution is representative of the distribution from which future points will be drawn is perhaps more often incorrect than correct” [69]. Even if a design distribution is initially representative, as noted earlier, representativeness is time- and place- dependent. This fact becomes all the more important in ML given (1) the tendency for ML-driven systems to be applied across contexts, driven in part by the financial advantages of scalability [144] and (2) the fact that these systems are implemented in dynamic environments where, “in more extreme cases, actions influenced by a model may alter the environment, invalidating future predictions” [98]. These tendencies are also part of how we critique these systems, warning of “portability trap[s]” [139] and context-dependence [144], of “zombie predictions” unable to account for policy change [87], or of “lack of representativeness” as a source of injustice when risk tools developed on one population are used to make decisions about another [68].

While not all portability challenges over time and space can be detected in terms of “shifts in the joint distribution of features and labels” [139], a growing body of work seeks to address such deviations from the “representativeness” or “matching distribution” assumption in machine learning. When only a few “out of distribution” data points are at play, practitioners speak more of outlier or anomaly detection, but when those points reflect meaningful changes in the underlying data-generating distribution, the problem is conceptualized as “shift” [88, 131, 141, 150]. The goals of this work are to detect and model various forms of shift for the sake of adjustment or at least “failing loudly” when a model receives inputs that do not come from the distribution on which it was trained [131]. Though translating that theoretical work to actively detecting and responding to shift in commercial, bureaucratic, and policy contexts remains a challenge, shift work is motivated by an empirical question: do distributions match?

Yet as we saw in African Americans’ struggles against discrimination in life insurance, the prospect of *the future* throws another wrench in this picture. Of course, if we are to predict at all, we must assume some link between data and the unobservable future, but that link raises normative dilemmas. Again, the world may not be as we wish it to be and people may push for desirable and aspirational outcomes not yet reflected in data [37, 83, 86]. “We” is a bit deceptive here: people may disagree on what those outcomes should be or on the morality of predictive practices [85]. These disagreements have high stakes, for their outcomes are performative—that is, acting on

data today (or not), even when they are currently, by some standard, accurate and representative, shapes future possibilities.

5.3 Participation

Representativeness finds further expression in calls for participatory processes intended to ensure equitable and inclusive ML-driven decision-making. A growing area of work urges greater opportunity for stakeholders—especially members of affected communities—to shape ML systems [95, 105, 144]. For example, the call for papers for the 2020 ICML Workshop on Participatory Approaches to Machine Learning decried the fact that “the design choices of a few” often determine how a system operates and espoused “more democratic, cooperative, and participatory” processes that “encourage the perspectives of those impacted by an ML system” [159]. Lee *et al.* design a framework to “promot[e] representative participation” in order to “empower stakeholders who typically do not have a say in the algorithms that govern their services, communities, or organizations” [95]. Sloane *et al.* describe “representative and consultative participation” as a means of obtaining the “right information from the right people,” though they argue such participation is not a panacea for ML’s ills [144]. Here, we again see representativeness linked to both political and epistemic aims.

One such aim, as we discussed in section 4.1, is to ensure the right to have a meaningful voice in designing systems that affect one’s community. Representativeness in this sense is a matter of design justice, as exemplified in the phrase “nothing about us without us” [41]. Representative processes are further espoused as an epistemic means to more effective and equitable systems, based on the prospect that community stakeholders hold substantial information about how systems are likely to work in practice and what impacts they will have on affected communities [10]. Without adequate representation of this knowledge in deliberative processes about ML systems, these impacts are likely to be underappreciated. Representativeness may also lend democratic legitimacy to ML systems that are tailored to community goals, rather than simply imposed “out of the box” without regard for local conditions [95].

Even if we decide that a decision-making process *should* be representative, we must still determine *of whom* it should be representative. Defining the population relative to which we assess representativeness is a crucial decision but not an obvious one. When we espouse representative or participatory processes in ML decision-making, we typically are not aiming for *national* representativeness; rather, we often have in mind some sort of *targeted* representativeness of an “affected population”—that is, “populations who have direct experience with the public system in question” [32]—or a collection of “stakeholders,” a term which is itself contested [112]. A concomitant concern is determining what attributes are important to represent—as Borges’ allegory illustrates, this requires judgment and line-drawing. Moreover, as in data collection, even when there is a demand for participants to be “representative” of some population, it is not always clear whether this means simply the inclusion of a variety of perspectives, representation *proportionate* to some population, or the *over*-representation of historically marginalized views in order to amplify them [102].⁶

⁶What’s more, these processes have an active and ongoing quality. Proponents of participatory ML speak of having voices “at the table” *throughout* decision-making

Representativeness, for all its virtues, might not achieve other normative goals. Many of the most fundamental concerns raised by algorithmic processes involve key civil and constitutional rights, like equal protection, freedom from discrimination, and due process [134]. Likeness-based matching—the idea that decision-makers should share attributes with the population, so as to form a “miniature” of the community (see sec. 4.1)—might be an ineffective prophylactic against the denial of these rights. Specialized expertise necessary to ensure these rights may come from institutional affiliates (like legal advocates) who work *with* affected groups but are not necessarily themselves members of it [102]. In this sense, some participatory processes echo the logic of early crop reporting—in which reporters were not intended to match the farmer population, but were valued as representatives for their specialized knowledge. This is not to say that expert advocates alone provide adequate representation. As Costanza-Chock demonstrates, strategies in which designers use techniques to “stand in for” engagement with the community itself risk embedding invalid assumptions about what its members want [41]. Participatory ML must consider how to balance representativeness alongside other logics and values.

Finally, participatory ML raises the key question of whether representation in a decision-making process translates to a system that benefits and meaningfully distributes power to those represented. To the degree that representativeness bestows legitimacy on an ML decision-making process, the term may be marshaled for rhetorical purposes—much as we saw in section 4.3. Just as proclaiming Gallup polls “representative” could paper over methodological details and questions about who most benefited from them, ML “participation-washing” might confer legitimacy on a decision-making process even when that process tokenizes, extracts value from, or further disenfranchises the people ostensibly represented. Debates about representativeness might also distract attention from whether the technology should be deployed at all [41, 74, 144].

5.4 Power

We close on the issue of data and power. Representativeness is a concept and value situated in social power dynamics. Despite the positive connotation it gains from its links to representation and inclusion, the apparent inclusiveness of representative data is not a guarantee of ethical data collection or control by data subjects over how data are used, nor does it ensure that those represented in a dataset will benefit from its use. Faced with a problem of unrepresentativeness, collecting more data can be an appealing solution for governments and companies that derive power, knowledge, and value from data collection—power not only over those they directly observe in their data, but also, via statistical inference, over the populations represented by those data [14, 85]. Such data collection may benefit that population, but the pursuit of representative data may also lead to exploitation and surveillance—particularly of those least able to resist [59, 73, 74, 129]. In fact, a narrow focus on collecting more data to improve a given technology could overshadow the issue of whether that data collection or the technology itself is desirable in the first place [8, 67]. The question is: representativeness for whom? For whose benefit or harm?

processes—including system formulation, deployment, and post-hoc evaluation. We do not only “do representativeness” at a moment in time and consider the matter settled.

The potential for data to be used for harm (even or especially when they are accurate) has driven some to call not for inclusion, better coverage, or representativeness, but for obfuscation and non-collection [17, 33, 72]—or as a 2019 workshop on these issues aptly framed it, “Please Don’t Include Us” [94]. The comparative harms of inclusion and exclusion from datasets are complex, contested and contextual. We do not aim to resolve them here, and they are thoughtfully considered by other scholars [20, 74]; but for our purposes, it is crucial to note how these tensions complicate the positive connotations often associated with representativeness.

For all the difficulties of making ethical judgements about inclusion and exclusion *within* the datasets we collect, questions of power also attend those we do not collect. This is, in a sense, the converse of the “model organism” datasets we discussed in 5.1; while some datasets are used widely across domains, others are not collected at all. Mimi Onuoha’s “Library of Missing Datasets” [123] and the Centre for Public Numbers’ “Missing Numbers” project [57] highlight the omission of certain datasets (e.g. statistics on police brutality) from otherwise data-rich environments. As Onuoha notes, missing datasets are not accidental, but may occur because powerful actors are not inclined to collect them. In response, efforts to collect *counterdata* [49] seek to fill these gaps and document phenomena that are important to marginalized groups. Sometimes, this involves not collecting more data about *those* groups, but “studying up,” a notion originating in anthropology [117] and recently raised in FAccT [13]. The call to “study up” implores us to reorient our research questions—and, concomitantly, our datasets—by making the powerful the targets of analysis. These efforts are oriented toward a sort of field-level representativeness, aiming to ensure that the field neither underscrutinizes the powerful nor underrepresents the interests of the less powerful.

6 CONCLUSION

In exploring representativeness, our challenge has been to open the concept’s ambiguities in a useful way without collapsing them to a single definition. We find that representativeness, with its roots in statistical sampling, is often a vaguely technical description linked alternatively to the typical example, random sample, proportionate miniature, and to the basic aims of scientific inference. As a political concept, it is used to communicate value for likeness and inclusion—but at times to exclude. Its combined scientific status and political resonance give it rhetorical authority, merited or not. For both data and politics, it is complicated by the roles power and aspiration play in shaping our futures. FAccT brings many of these dynamics together. But in exploring representativeness, we have also teetered at the edge of a rabbit hole, relying on additional suitcase words like fairness, normativity, and power. Hazardous as this may be, we could not do otherwise. The web of concepts that surround the concept of representativeness complicates it—but also gives it its resonance. Again: representativeness can rarely stand alone.

ACKNOWLEDGMENTS

We thank Kim Weeden, Marty Wells, Malte Ziewitz, and Emmanuel Didier, as well as members of Cornell’s AI, Policy, and Practice Initiative and College Scholar Program, the John D. and Catherine T. MacArthur Foundation, and the Miami Foundation.

REFERENCES

- [1] 1929. Adventurous Life Plea Stirs Miamians' Comment. *The Miami News* (July 28 1929).
- [2] 1929. Some Statistics on Main Street. *The Daily Republican* (Jan 22 1929).
- [3] 1929. A Summons to the Adventurous Life. *The Miami News* (July 21 1929).
- [4] 1929. The Times Book Table. *The Montclair Times* (Jan 26 1929).
- [5] 1930. Lester Leake Riley to Speak at Forum Sunday. *The Chatham Press* (Mar 22 1930).
- [6] 1930. 'Truth About the American Town is Pictured in Middletown,' says Mrs. Croxton to Sorosis Study Club. *St. Cloud Times* (Feb 19 1930).
- [7] Joshua T. Abbott, Katherine A. Heller, Zoubin Ghahramani, and Thomas L. Griffiths. 2011. Testing a Bayesian Measure of Representativeness Using a Large Image Database. In *Proceedings of the 24th International Conference on Neural Information Processing Systems* (Granada, Spain) (NIPS'11). Curran Associates Inc., Red Hook, NY, USA, 2321–2329.
- [8] 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*. 252–260.
- [9] Tanja Aitamurto, Jorge Saldivar Galli, and Juho Salminen. 2014. Self-selection in crowdsourced democracy: A bug or a feature?. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*, Baltimore, MD.
- [10] Tanja Aitamurto and Hélène Landemore. 2016. Crowdsourced Deliberation: The Case of the Law on Off-Road Traffic in Finland. *Policy & Internet* 8, 2 (2016), 174–196. <https://doi.org/10.1002/poi3.115>
- [11] Kendra Albert, Maggie Delano, Jonathon Penney, Afsaneh Rigot, and Ram Shankar Siva Kumar. 2020. Ethical Testing in the Real World: Evaluating Physical Testing of Adversarial Machine Learning. arXiv:2012.02048 [cs.CV]
- [12] Peter M. Aronow and Cyrus Samii. 2016. Does Regression Produce Representative Estimates of Causal Effects? *American Journal of Political Science* 60, 1 (2016), 250–267. <https://doi.org/10.1111/ajps.12185> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajps.12185>
- [13] 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*. 167–176.
- [14] Solon Barocas and Karen Levy. 2020. Privacy Dependencies. *Wash. L. Rev.* 95 (2020), 555.
- [15] Howard S Becker. 1996. The epistemology of qualitative research. *Ethnography and human development: Context and meaning in social inquiry* 27 (1996), 53–71.
- [16] Joseph A. Becker and C. L. Harlan. 1939. Developments in Crop and Livestock Reporting since 1920. *Journal of Farm Economics* 21, 4 (1939), 799–827. <http://www.jstor.org/stable/1231786>
- [17] Alvaro Bedoya. 2014. Big data and the underground railroad. *Slate* (2014). http://www.slate.com/articles/technology/future_tense/2014/11/big_data_underground_railroad_history_says_unfettered_collection_of_data.html
- [18] D.R. Bellhouse. 1988. A brief history of random sampling methods. *Handbook of Statistics* 6 (1988), 1–14. [https://doi.org/10.1016/S0169-7161\(88\)06003-1](https://doi.org/10.1016/S0169-7161(88)06003-1)
- [19] James R Beniger. 1983. Comments. *The Public Opinion Quarterly* 47, 4 (1983), 479.
- [20] Ruha Benjamin. 2019. *Race After Technology: Abolitionist tools for the new jim code*.
- [21] Jessica A. Bolker. 1995. Model systems in developmental biology. *BioEssays* 17, 5 (1995), 451–455. <https://doi.org/10.1002/bies.950170513>
- [22] Jorge Luis Borges. 1977. *The book of sand*. Dutton.
- [23] Tomas Borovicka, Marcel Jirina Jr., Pavel Kordik, and Marcel Jirina. 2012. Selecting Representative Data Sets. In *Advances in Data Mining Knowledge Discovery and Applications*, Adem Karahoca (Ed.). IntechOpen, Rijeka, Chapter 2. <https://doi.org/10.5772/50787>
- [24] Dan Bouk. 2015. *How Our Days Became Numbered: Risk and the Rise of the Statistical Individual*. University of Chicago Press, Chicago, IL.
- [25] Geoffrey C Bowker and Susan Leigh Star. 1999. *Sorting Things Out: Classification and Its Consequences*. MIT Press, Cambridge, MA.
- [26] Arthur L. Bowley. 1906. Address to the Economic Science and Statistics Section of the British Association for the Advancement of Science, York, 1906. *Journal of the Royal Statistical Society* 69, 3 (1906), 540–558. <http://www.jstor.org/stable/2339344>
- [27] Arthur L. Bowley. 1910. *An Elementary Manual of Statistics*. Macdonald and Evans, London, UK. <http://www.archive.org/details/cu31924013702950>
- [28] Arthur L. Bowley. 1926. Memorandum on Measurement of the Precision Attained in Sampling. *Bulletin de l'Institut international de statistique* 22, 1 (1926), 1–62.
- [29] Arthur L. Bowley and Alexander R. Burnett-Hurst. 1915. *Livelihood and Poverty: a Study in the Economic Conditions of Working-Class Households in Northampton, Warrington, Stanley and Reading*. G. Bell and Sons, London, UK.
- [30] Danah Boyd and Kate Crawford. 2012. Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, communication & society* 15, 5 (2012), 662–679.
- [31] Annie Britton, Martin McKee, Nick Black, Klim McPherson, Colin Sanderson, and Chris Bain. 1999. Threats to Applicability of Randomised Trials: Exclusions and Selective Participation. *Journal of Health Services Research & Policy* 4, 2 (1999), 112–121. <https://doi.org/10.1177/135581969900400210> PMID: 10387403.
- [32] Anna Brown, Alexandra Chouldechova, Emily Putnam-Hornstein, Andrew Tobin, and Rhema Vaithianathan. 2019. Toward algorithmic accountability in public services. In *CHI Conference on Human Factors in Computing Systems*, Glasgow, Scotland. <https://doi.org/10.1145/3290605.3300271>.
- [33] Finn Brunton and Helen Nissenbaum. 2015. *Obfuscation: A user's guide for privacy and protest*. MIT Press.
- [34] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. 77–91.
- [35] Danilo Bzdok, Naomi Altman, and Martin Krzywinski. 2018. Statistics versus Machine Learning. *Nature Methods* 15, 4 (2018), 233–34. <https://doi.org/10.1038/nmeth.4642>
- [36] Ernest Callenbach and Michael Phillips. 1985. *A Citizen Legislature*. Bayan Tree Books/Clear Glass, Berkeley/Bodega, CA.
- [37] Elisa Celis, Vijay Keswani, Damian Straszak, Amit Deshpande, Tarun Kathuria, and Nisheeth Vishnoi. 2018. Fair and Diverse DPP-Based Data Summarization (*Proceedings of Machine Learning Research*, Vol. 80). Jennifer Dy and Andreas Krause (Eds.). PMLR, Stockholmsmässan, Stockholm Sweden, 716–725. <http://proceedings.mlr.press/v80/celis18a.html>
- [38] Alexandra Chouldechova and Aaron Roth. 2018. The frontiers of fairness in machine learning. *arXiv preprint arXiv:1810.08810* (2018).
- [39] M.D. Christodoulou, J.A. Brettschneider, and D. Steinsaltz. 2020. Erosion of representativeness in a cohort study. *medRxiv* (2020). <https://www.medrxiv.org/content/early/2020/11/09/2020.02.13.20022012>
- [40] R. H. Coats. 1931. Enumeration and Sampling in the Field of the Census. *J. Amer. Statist. Assoc.* 26, 175 (1931), 270–284. <http://www.jstor.org/stable/2277905>
- [41] Sasha Costanza-Chock. 2020. *Design justice: Community-led practices to build the worlds we need*. MIT Press.
- [42] Mick P. Couper. 2013. Is the Sky Falling? New Technology, Changing Media, and the Future of Surveys. *Survey Research Methods* 7, 3 (Dec. 2013), 145–156. <https://doi.org/10.18148/srm/2013.v7i3.5751>
- [43] Lorraine Daston and Peter Galison. 1992. The Image of Objectivity. *Representations* 40 (1992), 81–128. <https://doi.org/doi:10.2307/2928741>
- [44] A. de P.R. Camargo. 2009. Sociology of Statistics: possibilities of a new field of investigation. *História, Ciências, Saúde-Manguinhos* 16, 4 (2009), 903–925. <https://doi.org/doi:10.1590/s0104-59702009000400004>
- [45] Emily Denton, Alex Hanna, Razvan Amironesei, Andrew Smart, Hilary Nicole, and Morgan Klaus Scheuerman. 2020. Bringing the People Back In: Contesting Benchmark Machine Learning Datasets. *arXiv preprint arXiv:2007.07399* (2020).
- [46] Alain Desrosières. 1991. How to Make Things Which Hold Together: Social Science, Statistics and the State. In *Discourses on Society: The Shaping of the Social Science Disciplines*, Peter Wagner, Björn Wittrock, and Richard Whitley (Eds.). Springer Netherlands, Dordrecht, 195–218. https://doi.org/10.1007/978-0-585-29174-1_8
- [47] Alain Desrosières. 1998. *The Politics of Large Numbers: a History of Statistical Reasoning*. Harvard University Press, Cambridge, MA.
- [48] Emmanuel Didier. 2020. *America By the Numbers: Quantification, Democracy, and the Birth of National Statistics*. MIT Press, Cambridge, MA. Page numbers may be off as they are based on unpublished copy obtained from the author in 2019.
- [49] Catherine D'Ignazio and Lauren F Klein. 2020. *Data feminism*. MIT Press.
- [50] Shah Ebrahim and George Davey Smith. 2013. Commentary: Should we always deliberately be non-representative? *International Journal of Epidemiology* 42, 4 (08 2013), 1022–1026. <https://doi.org/10.1093/ije/dyt105>
- [51] George C. Edwards, Lawrence R. Jacobs, Robert Y. Shapiro, and Michael X. Delli Carpini. 2011. Constructing Public Opinion: A Brief History of Survey Research. In *The Oxford Handbook of American Public Opinion and the Media*. Oxford University Press.
- [52] J Mark Elwood. 2013. Commentary: On representativeness. *International Journal of Epidemiology* 42, 4 (08 2013), 1014–1015. <https://doi.org/10.1093/ije/dyt101>
- [53] Nathan Ensmenger. 2012. Is chess the drosophila of artificial intelligence? A social history of an algorithm. *Social Studies of Science* 42, 1 (2012), 5–30. <https://doi.org/10.1177/0306312711424596> PMID: 22530382.
- [54] Feng Pan, Wei Wang, A. K. H. Tung, and Jiong Yang. 2005. Finding representative set from massive data. In *Fifth IEEE International Conference on Data Mining (ICDM'05)*.
- [55] William A. Firestone. 1993. Alternative Arguments for Generalizing From Data as Applied to Qualitative Research. *Educational Researcher* 22, 4 (1993), 16–23. <https://doi.org/10.3102/0013189X022004016>
- [56] Jill A. Fisher and Corey A. Kalbaugh. 2011. Challenging Assumptions About Minority Participation in US Clinical Research. *American Journal of Public Health* 101, 12 (2011), 2217–2222. <https://doi.org/10.2105/AJPH.2011.300279> PMID: 22021285.

- [57] Centre for Public Data. [n.d.]. <https://missingnumbers.org/>
- [58] Archon Fung. 2003. Survey Article: Recipes for Public Spheres: Eight Institutional Design Choices and Their Consequences. *Journal of Political Philosophy* 11, 3 (2003), 338–367. <https://doi.org/10.1111/1467-9760.00181>
- [59] Sidney Fussell. [n.d.]. How an Attempt at Correcting Bias in Tech Goes Wrong. <https://www.theatlantic.com/technology/archive/2019/10/google-allegedly-used-homeless-train-pixel-phone/599668/>. *The Atlantic* ([n.d.]).
- [60] George Gallup. 1938. Gallup Explains 'Opinion Sampling'. *New York Times* (May 17 1938).
- [61] George Gallup and Saul Forbes Rae. 1940. *The Pulse of Democracy: the Public Opinion Poll and How It Works*. Simon and Schuster, New York.
- [62] Henry E. Garrett. 1942. The Representativeness of a Sample. *The American Journal of Psychology* 55, 4 (1942), 580–581. <http://www.jstor.org/stable/1417131>
- [63] Timnit Gebru and Emily Denton. 2020. Tutorial on Fairness Accountability Transparency and Ethics in Computer Vision (CVPR '20). Association for Computing Machinery, New York, NY, USA.
- [64] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2018. Datasheets for datasets. *arXiv preprint arXiv:1803.09010* (2018).
- [65] Gerd Gigerenzer, Zeno Swijtink, Theodore Porter, Lorraine Daston, John Beatty, and Lorenz Krüger. 1989. *The Empire of Chance: How Probability Changed Science and Everyday Life*. Cambridge University Press, Cambridge, England.
- [66] Emma Goldberg. 2020. Cancer Projects to Diversify Genetic Research Receive New Grants. *The New York Times* (September 11 2020).
- [67] Kevin D Haggerty. 2009. Methodology as a knife fight: The process, politics and paradox of evaluating surveillance. *Critical Criminology* 17, 4 (2009), 277.
- [68] Melissa Hamilton. 2014. Adventures in Risk: Predicting Violent and Sexual Recidivism in Sentencing Law. *Arizona State Law Journal* (2014). <http://ssrn.com/abstract=2416918>
- [69] David J. Hand. 2006. Classifier Technology and the Illusion of Progress. *Statist. Sci.* 21, 1 (02 2006), 1–14. <https://doi.org/10.1214/088342306000000060>
- [70] Alex Hanna, Emily Denton, Andrew Smart, and Jamila Smith-Loud. 2020. Towards a Critical Race Methodology in Algorithmic Fairness (FAT* '20). Association for Computing Machinery, New York, NY, USA, 501–512. <https://doi.org/10.1145/3351095.3372826>
- [71] Sandra Harding. 1992. Rethinking standpoint epistemology: What is "strong objectivity"? *The Centennial Review* 36, 3 (1992), 437–470.
- [72] Nabil Hassein. 2017. Against black inclusion in facial recognition. <https://digitaltalkingdrum.com/2017/08/15/against-black-inclusion-in-facial-recognition/>
- [73] Amy Hawkins. 2018. Beijing's big brother tech needs African faces. *Foreign Policy* 24 (2018).
- [74] Anna Lauren Hoffmann. 2020. Terms of Inclusion: Data, Discourse, Violence. *New Media and Society* (2020).
- [75] Margaret H. Hogg. 1930. Sources of Incomparability and Error in Employment-Unemployment Surveys. *J. Amer. Statist. Assoc.* 25, 171 (1930), 284–294. <http://www.jstor.org/stable/2278191>
- [76] Lily Hu and Issa Kohler-Hausmann. 2020. What's Sex Got to Do with Machine Learning?. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (FAT* '20). Association for Computing Machinery, New York, NY, USA, 513. <https://doi.org/10.1145/3351095.3375674>
- [77] Sarah E. Igo. 2007. *The Averaged American: Surveys, Citizens, and the Making of a Mass Public*. Harvard University Press, Cambridge, MA.
- [78] Charles Isbell. 2020. NeurIPS Keynote: You Can't Escape Hyperparameters and Latent Variables: Machine Learning as a Software Engineering Enterprise (NeurIPS '20). Association for Computing Machinery.
- [79] Adolph Jensen. 1926. Report on the Representative Method in Statistics. *Bulletin de l'Institut international de statistique* 22, 1 (1926), 359–438.
- [80] Adolph Jensen. 1928. Purposive Selection. *Journal of the Royal Statistical Society* 91, 4 (1928), 541–547. <http://www.jstor.org/stable/2341671>
- [81] Caradog Jones. 1910. *A First Course in Statistics*. G. Bell and Sons, London, UK. <http://www.archive.org/details/cu31924014485670>
- [82] Orange Judd. 1862. An Important Enterprise - The Co-operation of all our Readers Asked. *American Agriculturalist* 21 (mar 1862). <https://www.biodiversitylibrary.org/bibliography/144617>
- [83] Matthew Kay, Cynthia Matuszek, and Sean A. Munson. 2015. Unequal Representation and Gender Stereotypes in Image Search Results for Occupations. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 3819–3828. <https://doi.org/10.1145/2702123.2702520>
- [84] Anders Nicolai Kiaer. 1976. *Den repræsentative undersøgelsesmetode = The representative method of statistical surveys*. Oslo: Aschehoug.
- [85] Barbara Kiviat. 2019. The Moral Limits of Predictive Practices: The Case of Credit-Based Insurance Scores. *American Sociological Review* 84, 6 (2019), 1134–1158.
- [86] Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Cass R Sunstein. 2018. Discrimination in the Age of Algorithms. *Journal of Legal Analysis* 10 (2018).
- [87] John Logan Koepke and David G Robinson. 2018. Danger ahead: Risk assessment and the future of bail reform. *Wash. L. Rev.* 93 (2018), 1725.
- [88] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Sara Beery, et al. 2020. WILDS: A Benchmark of in-the-Wild Distribution Shifts. *arXiv preprint arXiv:2012.07421* (2020).
- [89] William Kruskal. 1978. LANGUAGE: Formulas, Numbers, Words: Statistics in Prose. *The American Scholar* 47, 2 (1978), 223–229. <http://www.jstor.org/stable/41210378>
- [90] William Kruskal and Frederick Mosteller. 1979. Representative Sampling, I: Non-Scientific Literature. *International Statistical Review / Revue Internationale de Statistique* 47, 1 (1979), 13–24. <http://www.jstor.org/stable/1403202>
- [91] William Kruskal and Frederick Mosteller. 1979. Representative Sampling, II: Scientific Literature, Excluding Statistics. *International Statistical Review / Revue Internationale de Statistique* 47, 2 (1979), 111–127. <http://www.jstor.org/stable/1402564>
- [92] William Kruskal and Frederick Mosteller. 1979. Representative Sampling, III: The Current Statistical Literature. *International Statistical Review / Revue Internationale de Statistique* 47, 3 (1979), 245–265. <http://www.jstor.org/stable/1402647>
- [93] William Kruskal and Frederick Mosteller. 1980. Representative Sampling, IV: The History of the Concept in Statistics, 1895–1939. *International Statistical Review / Revue Internationale de Statistique* 48, 2 (1980), 169–195. <http://www.jstor.org/stable/1403151>
- [94] Digital Justice Lab. [n.d.]. Call for Participation: Please Don't Include Us. <https://digitaljusticelab.ca/cfp>. Accessed: 2020-09-30.
- [95] Min Kyung Lee, Daniel Kusbit, Anson Kahng, Ji Tae Kim, Xinran Yuan, Allissa Chan, Daniel See, Ritesh Noothigattu, Siheon Lee, Alexandros Psomas, et al. 2019. WeBuildAI: Participatory framework for algorithmic governance. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–35.
- [96] Jonas Lerman. 2013. Big data and its exclusions. *Stan. L. Rev. Online* 66 (2013), 55.
- [97] C. S. Lewis. 1960. *Studies in Words* (1st. ed.). Cambridge University Press, Cambridge, England.
- [98] Zachary C Lipton. 2018. The mythos of model interpretability. *Queue* 16, 3 (2018), 31–57.
- [99] Zachary C Lipton and Jacob Steinhardt. 2018. Troubling trends in machine learning scholarship. *arXiv preprint arXiv:1807.03341* (2018).
- [100] Sharon L. Lohr. 1999. *Sampling: Design and Analysis*. Duxbury Press, Pacific Grove, CA.
- [101] Robert Lynd and Helen Lynd. 1929. *Middletown: A Study in Contemporary American Culture*. Harcourt, Brace and Company, New York.
- [102] Lassana Magassa, Meg Young, and Batya Friedman. 2017. *Diverse Voices: A How-To Guide for Facilitating Inclusiveness in Tech Policy*. Technical Report. University of Washington Tech Policy Lab. https://techpolicylab.uw.edu/wp-content/uploads/2017/10/TPL_Diverse_Voices_How-To_Guide_2017.pdf
- [103] J. G. Mandello. 1905. The Future of Statistics. *Journal of the Royal Statistical Society* 68, 4 (1905), 725–732. <http://www.jstor.org/stable/2339425>
- [104] Bernard Manin. 1997. *The Principles of Representative Government*. Cambridge University Press, New York.
- [105] Donald Martin Jr, Vinod Prabhakaran, Jill Kuhlberg, Andrew Smart, and William S Isaac. 2020. Participatory Problem Formulation for Fairer Machine Learning Through Community Based System Dynamics. *arXiv preprint arXiv:2005.07572* (2020).
- [106] Viktor Mayer-Schönberger and Kenneth Cukier. 2013. *Big Data: a Revolution That Will Transform How We Live, Work, and Think*. Houghton Mifflin Harcourt, Boston, MA.
- [107] Quinn McNemar. 1940. Sampling in psychological research. *Psychological Bulletin* 37, 6 (1940), 331–365. <https://doi.org/10.1037/h0063476>
- [108] C B Mervis and E Rosch. 1981. Categorization of Natural Objects. *Annual Review of Psychology* 32, 1 (1981), 89–115. <https://doi.org/10.1146/annurev.ps.32.020181.000513>
- [109] Martine Mespoulet. 2002. From Typical Areas to Random Sampling: Sampling Methods in Russia from 1875 to 1930. *Science in Context* 15, 3 (2002), 411–425. <https://doi.org/10.1017/S0269889702000546>
- [110] Susan M Miller and Lael R Keiser. 2020. Representative Bureaucracy and Attitudes Toward Automated Decision Making. *Journal of Public Administration Research and Theory* (2020).
- [111] Marvin Minsky. 2006. *The Emotion Machine*. Simon & Schuster, New York.
- [112] Ronald K Mitchell, Bradley R Agle, and Donna J Wood. 1997. Toward a theory of stakeholder identification and salience: Defining the principle of who and what really counts. *Academy of management review* 22, 4 (1997), 853–886.
- [113] Shira Mitchell, Eric Potash, Solon Barocas, Alexander D'Amour, and Kristian Lum. 2018. Prediction-based decisions and fairness: A catalogue of choices, assumptions, and definitions. *arXiv preprint arXiv:1811.07867* (2018).
- [114] Mara Mordecai. 2020. *How Americans envision a post-pandemic world order*. Technical Report. Pew Research. <https://www.pewresearch.org/fact-tank/2020/06/02/how-americans-envision-a-post-pandemic-world-order/>

- [115] Fred Morstatter, Jürgen Pfeffer, and Huan Liu. 2014. When is It Biased? Assessing the Representativeness of Twitter's Streaming API. In *Proceedings of the 23rd International Conference on World Wide Web (Seoul, Korea) (WWW '14 Companion)*. Association for Computing Machinery, New York, NY, USA, 555–556. <https://doi.org/10.1145/2567948.2576952>
- [116] Deirdre K. Mulligan, Joshua A. Kroll, Nitin Kohli, and Richmond Y. Wong. 2019. This Thing Called Fairness: Disciplinary Confusion Realizing a Value in Technology. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 119 (Nov. 2019). <https://doi.org/10.1145/3359221>
- [117] Laura Nader. 1972. Up the anthropologist: perspectives gained from studying up. (1972).
- [118] Arvind Narayanan. 2018. Translation tutorial: 21 fairness definitions and their politics. In *Proc. Conf. Fairness Accountability Transp., New York, USA*, Vol. 1170.
- [119] Jerzy Neyman. 1934. On the Two Different Aspects of the Representative Method: The Method of Stratified Sampling and the Method of Purposive Selection. *Journal of the Royal Statistical Society* 97, 4 (1934), 558–625. <http://www.jstor.org/stable/2342192>
- [120] Ellen A Nohr and Jørn Olsen. 2013. Commentary: Epidemiologists have debated representativeness for more than 40 years—has the time come to move on? *International Journal of Epidemiology* 42, 4 (08 2013), 1016–1017. <https://doi.org/10.1093/ije/dyt102>
- [121] Sofia C. Olhede and Patrick J. Wolfe. 2018. The future of statistics and data science. *Statistics & Probability Letters* 136 (2018), 46–50. <https://doi.org/10.1016/j.spl.2018.02.042>
- [122] John O'Neill. 2001. Representing People, Representing Nature, Representing the World. *Environment and Planning C: Government and Policy* 19, 4 (2001), 483–500. <https://doi.org/10.1068/c12s>
- [123] Mimi Onuoha. [n.d.]. <http://mimionuoha.com/the-library-of-missing-datasets>
- [124] Amandalynne Paullada, Inioluwa Deborah Raji, Emily M Bender, Emily Denton, and Alex Hanna. 2020. Data and its (dis)contents: A survey of dataset development and use in machine learning research. *arXiv preprint arXiv:2012.05345* (2020).
- [125] Anne Phillips. 1995. *The Politics of Presence*. Clarendon Press, Oxford, UK.
- [126] Hanna Fenichel Pitkin. 1972. *The Concept of Representation*. University of California Press, Berkeley, NY.
- [127] Theodore M. Porter. 1986. *The Rise of Statistical Thinking, 1820-1900*. Princeton University Press, Princeton, NJ.
- [128] Theodore M Porter. 1996. *Trust in numbers: The pursuit of objectivity in science and public life*. Princeton University Press.
- [129] Julia Powles and Helen Nissenbaum. 2018. The seductive diversion of 'solving' bias in artificial intelligence. (2018).
- [130] Jean-Guy Prévost and Jean-Pierre Beaud. 2012. *Statistics, Public Debate and the State, 1800-1945: a Social, Political and Intellectual History of Numbers*. Pickering & Chatto, London, UK.
- [131] Stephan Rabanser, Stephan Günnemann, and Zachary Lipton. 2019. Failing loudly: An empirical study of methods for detecting dataset shift. In *Advances in Neural Information Processing Systems*. 1396–1408.
- [132] Inioluwa Deborah Raji and Joy Buolamwini. 2019. Actionable auditing: Investigating the impact of publicly naming biased performance results of commercial ai products. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 429–435.
- [133] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why should I trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 1135–1144.
- [134] R Richardson, JM Schultz, and VM Southerland. 2019. Litigating Algorithms 2019 Report: New Challenges to Government Use of Algorithmic Decision Systems.
- [135] Lorenzo Richiardi, Costanza Pizzi, and Neil Pearce. 2013. Commentary: Representativeness is usually not necessary and often should be avoided. *International Journal of Epidemiology* 42, 4 (08 2013), 1018–1022. <https://doi.org/10.1093/ije/dyt103>
- [136] Kenneth J Rothman, John EJ Gallacher, and Elizabeth E Hatch. 2013. Why representativeness should be avoided. *International Journal of Epidemiology* 42, 4 (08 2013), 1012–1014. <https://doi.org/10.1093/ije/dys223>
- [137] Matthew J Salganik. 2019. *Bit by bit: Social research in the digital age*. Princeton University Press.
- [138] Libby Schweber. 2006. *Disciplining Statistics: Demography and Vital Statistics in France and England 1830/1885*. Duke University Press, Durham, NC. <https://doi.org/10.1215/9780822388524>
- [139] Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 59–68.
- [140] You Poh Seng. 1951. Historical Survey of the Development of Sampling Theories and Practice. *Journal of the Royal Statistical Society. Series A (General)* 114, 2 (1951), 214–231. <http://www.jstor.org/stable/2980977>
- [141] A Shafaei, M Schmidt, and JJ Little. 2018. Does your model know the digit 6 is not a cat. A Less Biased Evaluation of "Outlier" Detectors. *ArXiv e-Print* (2018).
- [142] Shreya Shankar, Yoni Halpern, Eric Breck, James Atwood, Jimbo Wilson, and D. Sculley. 2017. No Classification without Representation: Assessing Geodiversity Issues in Open Data Sets for the Developing World. (2017). arXiv:1711.08536 [stat.ML]
- [143] Giorgio Sirugo, Scott M. Williams, and Sarah A. Tishkoff. 2019. The Missing Diversity in Human Genetic Studies. *Cell* 177, 1 (2019), 26–31. <https://doi.org/10.1016/j.cell.2019.02.048>
- [144] Mona Sloane, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 2020. Participation Is Not a Design Fix for Machine Learning. In *Proceedings of the 37th International Conference on Machine Learning*. Vienna, Austria.
- [145] Mario Luis Small. 2009. How many cases do I need? On science and the logic of case selection in field-based research. *Ethnography* 10, 1 (2009), 5–38.
- [146] Isabel Stamm and Lina Eklund. 2017. With Great Power Comes Great Responsibility: Crowdsourcing Raises Methodological and Ethical Questions for Academia. *Impact of Social Sciences (blog)* (5 apr 2017). <https://blogs.lse.ac.uk/impactofsocialsciences/2017/04/05/crowdsourcing-raises-methodological-and-ethical-questions-for-academia/>
- [147] Frederick F. Stephan. 1939. Representative Sampling in Large-Scale Surveys. *J. Amer. Statist. Assoc.* 34, 206 (1939), 343–352. <http://www.jstor.org/stable/2278856>
- [148] Frederick F. Stephan. 1941. Stratification in Representative Sampling. *Journal of Marketing* 6, 1 (1941), 38–46. <http://www.jstor.org/stable/1245763>
- [149] Frederick F. Stephan. 1948. History of the Uses of Modern Sampling Procedures. *J. Amer. Statist. Assoc.* 43, 241 (1948), 12–39. <http://www.jstor.org/stable/2280064>
- [150] Amos Storkey. 2009. When Training and Test Sets are Different: Characterizing Learning Transfer. In *Dataset Shift in Machine Learning*. Joaquin Quiñero-Candela (Ed.). MIT Press, Cambridge, Mass, 3–28.
- [151] Harini Suresh and John V. Gutttag. 2019. A Framework for Understanding Unintended Consequences of Machine Learning. *CoRR* abs/1901.10002 (2019). arXiv:1901.10002 <http://arxiv.org/abs/1901.10002>
- [152] Joshua B. Tenenbaum and Thomas L. Griffiths. 2001. The Rational Basis of Representativeness. In *23rd Annual Conference of the Cognitive Science Society*.
- [153] Zeynep Tufekci. 2014. Big questions for social media big data: representativeness, validity and other methodological pitfalls. In *Eighth International AAAI Conference on Weblogs and Social Media*.
- [154] Amos Tversky and Daniel Kahneman. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science* 185, 4157 (1974), 1124–1131. <http://www.jstor.org/stable/1738360>
- [155] Amos Tversky and Danie Kahneman. 1981. Judgments of and by Representativeness. (1981), 2–27. <https://apps.dtic.mil/docs/citations/ADA099502>
- [156] Sidney Verba, Kay Lehman Schlozman, and Henry E Brady. 1995. *Voice and Equality*. Harvard University Press, Cambridge, MA.
- [157] Raymond Williams. 1976. *Keywords: a Vocabulary of Culture and Society*. Oxford University Press, New York, NY.
- [158] 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 (FAT* '19). Association for Computing Machinery, New York, NY, USA, 191–200. <https://doi.org/10.1145/3287560.3287577>
- [159] Angela Zhou, David Madras, Inioluwa Deborah Raji, Bogdan Kulynych, Smitha Milli, and Richard Zemel. [n.d.]. Call for Participation: Participatory Approaches to Machine Learning. <https://participatoryml.github.io/>. Accessed: 2020-09-30.