

---

# Real Risks of Fake Data: Synthetic Data, Diversity-Washing and Consent Circumvention

---

Cedric Whitney<sup>1</sup> Justin Norman<sup>1</sup>

## Abstract

Machine learning systems require representations of the real world for training and testing - they require data, and lots of it. Collecting data at scale has logistical and ethical challenges, and synthetic data promises a solution to these challenges. Instead of needing to collect photos of real people's faces to train a facial recognition system, a model creator could create and use photo-realistic, synthetic faces. The comparative ease of generating this synthetic data rather than relying on collecting data has made it a common practice. We present two key risks of using synthetic data in model development. First, we detail the high risk of false confidence when using synthetic data to increase dataset diversity and representation. We base this in the examination of a real world use-case of synthetic data, where synthetic datasets were generated for an evaluation of facial recognition technology. Second, we examine how using synthetic data risks circumventing consent for data usage. We illustrate this by considering the importance of consent to the U.S. Federal Trade Commission's regulation of data collection and affected models. Finally, we discuss how these two risks exemplify how synthetic data complicates existing governance and ethical practice; by decoupling data from those it impacts, synthetic data is prone to consolidating power away those most impacted by algorithmically-mediated harm.

## 1. Introduction

Facial recognition technology (FRT) has become commonplace, used from flight check-in at airports to police crowd-monitoring. Bias in FRT models has resulted in mis-

---

<sup>1</sup>School of Information, University of California, Berkeley, Berkeley, United States of America. Correspondence to: Cedric Whitney <cedricwhitney@berkeley.edu>.

identification and expanded surveillance, causing unjust incarceration and other discriminatory outcomes. Attempts to solve these issues by increasing the accuracy of FRT run headfirst into problems; for a machine learning-based computer vision system to be considered robust enough for a given real-world task, it must “generalize” to images that vary widely in quality and domain (image granularity, race, age, gender, background, head pose, hats, glasses, etc.). Datasets with this level of granular design and annotation, that are also large enough for use in deep learning, are nearly impossible to find due to logistical and ethical concerns. As a result, researchers have turned to synthetic data generation, where data is generated to resemble something without being a representation of an instance of it — a drawing of a generic face as compared to a photograph of a real person. Synthetic data has been used to augment existing datasets and create new datasets for better training and evaluation of FRT models. Logistical and ethical challenges to data collection exist outside of FRT, and synthetic data usage has become commonplace across machine learning, from computer vision to large language models. This paper examines two key risks of using synthetic data.

Synthetic data is fundamentally useful where real data is not fit to task, necessitating that synthetic data must be both similar enough to be meaningful, but different enough to mitigate the reason the real data is not usable (Jordan et al., 2022). Jordan et al. propose three attributes of synthetic data that must be met for it to function in lieu of real data: utility, fidelity, and privacy. This paper focuses on facial recognition because it clearly articulates the risks of synthetic data, inherently forcing trade-offs between these attributes. There is high difficulty in making a picture of a face private but still usable as training data (privacy vs. utility) — a face which has been obscured to the point where an identity could not be gleaned is less useful (Chamikara et al., 2020). Achieving fidelity in facial datasets, a measure of how well synthetic data matches the real world, is also saliently difficult in facial recognition use cases, as we examine below.

The first risk we focus on is the *high risk of false confidence* in the ability of synthetic datasets to mitigate bias in data distribution and representation. We demonstrate this through the real-world example of using synthetic data for a facial

recognition model evaluation. This paper was motivated by the realization of the under-explored risks of synthetic data while conducting the evaluation, and we present it both to provide an example of how synthetic data is used and to detail the concerns that conducting it made apparent. In brief, synthetic data offers a way of diversifying datasets, but diversity in real-world faces often follows from cultural practices that are qualitative and meaning-laden rather than quantitative. Creating a synthetic dataset or adding synthetic data to existing datasets in an attempt to diversify that dataset runs the risk of *diversity-washing* — appearing to resolve valid criticism regarding a dataset’s distribution and representation but in a way that is superficial. As a result, using synthetic data risks legitimizing technologies such as FRT despite potentially continuing to propagate bias by achieving false fidelity.

The second risk we examine is how using synthetic data risks *circumventing consent for data usage*, illustrating the impacts by considering the importance of consent to the U.S. Federal Trade Commission’s regulation of data collection and affected models. Synthetic data provides an avenue for model developers to side-step thorny issues around collecting large-scale representative facial datasets. Proper consent to data usage is foundational to the privacy enforcement tools that the FTC has used to require companies delete ML models trained on improperly collected data, a key regulatory hurdle to improper data collection and resulting harmful model deployment. Using synthetic data risks circumventing and obfuscating consent, thus complicating deterrence and enforcement.

This paper proceeds as follows: We begin in Section 2 by summarizing related prior work. We first examine work on datasets and representation, before discussing participation and consent and power over data and models. Finally, we discuss synthetic data and its use. We then proceed to the two titular risks of synthetic data that this paper focuses on — diversity-washing (Section 3) and the circumvention of consent (Section 4). We draw upon two real-world examples: a facial recognition evaluation task conducted using synthetic data, and the FTC’s enforcement actions against models trained on deceptively collected data to illustrate these risks. Finally, in Section 5 we expand upon how these two risks are examples of irresponsible use of synthetic data: consolidating power in the hands of model creators, and decoupling data from those it represents and those who are harmed by its improper use. It is our intention for this research to contribute to the field by presenting tangible examples and background for the challenges inherent in responsible use of synthetic data, thus laying foundations for further work and debate. We call for future work to examine the breadth and usage of synthetic data and to work towards both mitigating synthetic data’s risks and enabling its potential for participatory empowerment.

## 2. Related Work

In this section, first, we focus on the datasets that underpin machine learning systems, and detail how that work treats the specific issue of data distribution and representation in those datasets (Section 2.1). Next, we discuss prior work on participatory governance, consent and data privacy, and attempts to capture some power over dataset creation and usage by those most affected (Section 2.2). Finally, we provide a summary of work detailing what synthetic data is and how it is used (Section 2.3).

### 2.1. Datasets, Diversity and Representation

Datasets are the hidden infrastructure behind machine learning, most visible when the systems built on them break (Jackson, 2014). Models developed and dependent on these large datasets can lead to biased and harmful effects, with models used for bureaucratic categorization in particular having a long history of harm (Bullock, 2019; Alkhatib & Bernstein, 2019; Spade, 2015). The collection of data is then frequently the starting point for ML-disseminated discrimination and bias in domains such as hiring (Raghavan et al., 2020), advertising (Lee et al., 2019), pricing (Wu et al., 2022), the application of law, and government allocation of resources (Abebe & Goldner, 2018); as well as being vital for identifying and enforcing against discrimination (Andrus & Villeneuve, 2022). The stakes of responsible dataset development, then, are high, and we build on critical previous work (Paullada et al., 2021; Peng et al., 2021; Hutchinson et al., 2021) in focusing on the ways that dataset creators have significant impact on the harms that occur downstream via their development, usage and deployment (Khan & Hanna, 2022). More narrowly, we hope to bring focus to important risks present when synthetic data is used in the process of creating and using datasets in machine learning development.

Datasets used for facial recognition models, where the goal of the model is matching an image or video of a face to an identity, have received much scrutiny — specifically for violating privacy (Harvey & LaPlace, 2021; Birhane & Prabhu, 2021). An analysis by Crawford & Paglen (Crawford & Paglen, 2021) of ImageNet, a frequently-used large dataset, demonstrated active labeling of faces with offensive and derogatory classifications, and Prabhu & Birhane (Birhane & Prabhu, 2021) make the point that beyond obvious privacy harms such as blackmail, the creation of one of these datasets causes similar datasets to propagate. The recent identification of CSAM material in the popular LAION dataset is an example of this (Thiel, 2023). In the case of synthetic data, where the data is frequently derivative of previously collected data (as expanded upon below in Section 2.3), this then risks the continued propagation of non-consensual imagery. The above highlights the need

to focus beyond just the models. Much of the critical AI literature focuses on vital interventions to change model outputs that discriminate against protected classes. This work, instead, more closely follows work such as Buolamwini & Gebru’s “Gender Shades” (Buolamwini & Gebru, 2018) that is focused on the *data* which is fundamental to AI system development.

To understand the risks posed by using synthetic data to create and add to datasets used in machine learning development, it is necessary to understand the landscape of both machine learning development and the stakeholders impacted by it. We will use the taxonomy of dataset development stages and subjects presented by Khan & Hanna (Khan & Hanna, 2022). We lean on this taxonomy throughout the paper, finding it to be clear in its intention of “providing a common language for conversations across datasets” between practitioners, scholars and regulators. Starting from first principles, machine learning is not rules-based like traditional software development, but instead consists of a practitioner teaching a model to identify patterns in a dataset. To begin, Khan & Hanna assert, the *task* for which the model is being trained must be formulated and constrained. Next, *data* must be collected, meeting the constraints of what is necessary to train a model to achieve said task. That data collection is usually broad, requiring the data be *cleaned* before it is annotated. The *labels* attributed to data by the annotator are of vital importance to machine learning systems, as the systems are taught to identify those labels in their training data. After this point, model *training*, *valuation*, *implementation*, etc. may occur. Khan & Hannah define multiple *stakeholders* in the process of creating the datasets used for ML development: the curator, the data annotator, the data subject, the copyright holder, and the model subject. The *curator* is the entity responsible for dataset creation, while the *annotator* is (frequently outsourced (Gray & Suri, 2019)) responsible for annotating the dataset. The *data subject* is the person whose biometric information is present within the collected data, the *copyright holder* may hold exclusive rights over that data<sup>1</sup>, and the *model subject* is the person who is impacted by the decisions made by the model trained on the data. The last three categories are fluid, and can consist of the same person or of two or three distinct entities.

Creating a training dataset that is representative of model subjects is challenging, and when done poorly, results in inaccurate and frequently harmful outputs; here we consider how such harm might arise across different stages. In producing a representative dataset, different dimensions of identities are frequently missed. Datasets are frequently biased by necessity to meet a specific intended use of a model, but

<sup>1</sup>not all data is copyright protected, and even when it is, different legal regimes have different limitations and exceptions to exclusive intellectual property rights

when categories are socially constructed (such as race and gender), the observed, inferred data that dataset curators use to bound and constrain data collection (and that annotators must use to label) can clash with how the data subject self-identifies. Biased data representation is also a concern, with annotation reflecting social biases and stereotypes across gender, race, and more (Scheuerman et al., 2020). This can lead to rampant misrepresentation and miscategorization of both data and model subjects, producing forms of control. Annotation work is inherently an interpretive project, but results in data that is perceived to be ground truth (Bowman & Dahl, 2021). In reality, as shown by Recht et al. in work on testing the generalizability of ImageNet, when attempts to replicate annotation are made, different distributional properties for the same data emerge (Recht et al., 2019). Annotation is also frequently the cause of artifacts in datasets that allow for models to overfit to training data when solving a task, with many of the concerns arising from how human data annotators are instructed to label (Yang et al., 2020). Beyond concerns with annotation, the question of what data to collect when constructing a dataset to correctly answer a question is complex. The task of annotation itself presupposes that the question being asked of a model is one that can be answered — Aguera y Arcas et al. demonstrate how a model trained on ‘gaydar’ data was in reality labeled around stereotyped aesthetic traits, showcasing an example of labels being generated not because of any model-pertinent aspects of the data, but rather simply because the question was being asked (y Arcas, 2018).

Beyond the above challenges to creating a representative dataset, there are ethical issues that are raised by efforts to produce a dataset that represents a diverse community. Data collection requires infrastructure, and that infrastructure is frequently co-constitutive with surveillance infrastructure. Even when data collection is initiated in service of providing services to the most disenfranchised, rendering the members of those communities hyper-visible frequently serves to hurt those same communities, as decisions are made for them by others (Andrus & Villeneuve, 2022). These decisions can reinforce oppressive norms, such as visual gender binaries (Bivens, 2017; Hamidi et al., 2018), further delegitimizing disenfranchised groups in a clear example of administrative violence (Spade, 2015). Even when categorization schema of data subjects are correct, their use as prescriptive instead of explanatory can lead to attribution errors, co-opting classification in an oppressed group as a reason for that very oppression. Machine learning systems used to predict recidivism are a prime example (Christin, 2017), where factors like race, which make a group member more likely to be targeted for discrimination, are frequently used instead as a predictive factor when individuals are made model subjects.



## 2.2. Participation, Consent and Privacy

Participatory approaches are frequently fronted as a way of mitigating the harm that results from AI systems, both at the dataset level as described above, and in model training and deployment. These approaches focus on engaging the public, and build on policy approaches such as feedback sessions, public hearings and impact assessments (Glucker et al., 2013; Hügel & Davies, 2020). Participatory design methods in particular focus on co-design to incorporate user context, needs and values (Bratteteig & Wagner, 2016; Sloane, 2022; Iversen et al., 2012; Halloran et al., 2009), designing systems *with* those affected instead of *for* them. Recently, participatory AI work has explicitly focused on those for which AI most frequently exacerbates harm (Irgens et al., 2022; Robertson & Salehi, 2020). Patel et al. (Patel et al., 2021) draw from previous work, including Arnstein’s influential Ladder of Citizen Participation (Arnstein, 1969), to detail five levels of participation in data stewardship, including: 1) *informing* people about how their data is used through methods such as model cards, 2) *consulting* people through UX research and surveys, 3) *involving* people in data governance through panels and public deliberation, 4) *co-design* of data governance and consequent technologies through structures such as data trusts, and 5) enabling *decision-making* through citizen-led governance boards. We will return to Patel’s framework when discussing the risk of synthetic data circumventing consent (Section 4).

There is a wide — and growing! — diversity of participatory work in AI. Examples range from crowdsourcing impacts (Barnett & Diakopoulos, 2022; Diaz et al., 2022) and data labeling (Park et al., 2019) to eliciting preferences for dataset collecting and design decisions (Christiano et al., 2023). Peng et al. (Peng et al., 2021) recommend that dataset creators make ethically salient information clear and accessible while actively stewarding the dataset and its future use, and encourage retrospective study of datasets due to the difficulties in understanding issues at the beginning. Hutchinson et al. (Hutchinson et al., 2021) detail documentation requirements at each stage of the dataset development lifecycle, with different document types for each stage, and call for frameworks for transparency and accountability. These fall across the range of Patel et al.’s framework, (Patel et al., 2021) and critiques of these methods include characterization of it as ‘participation-washing’ (Gilman, 2022; Sloane et al., 2020), with Arnstein describing approaches such as public requests for comment as “tokenizing” and “inadequate in shifting power” (Arnstein, 1969). Sloane et al. (Sloane et al., 2020) argue that these approaches can function as unrecognized labor, and the line between tokenization and participation in cases such as crowdsourcing is quite blurry. Birhane et al. (Birhane et al., 2022) show examples of community inclusion in annotating datasets, improving documentation and increasing the utility of large

language models for under-served languages, and other examples include community organizations such as the Detroit URC3 which evaluates potential partnerships between community organizations and researchers to avoid exploitation (Corbett et al., 2023), and examples from Indigenous Data Sovereignty (Rainie et al., 2019). At a large scale however, there are still major hurdles. Groves et al. investigate the hurdle of making participatory approaches work in the commercial AI labs that are the primary site for AI research, and find that “corporate profit motive and concern around exploitation are at present functioning as significant barriers to the use of participatory methods in AI” (Groves et al., 2023).

While the above participatory approaches center shifting decision-making power to include data subjects and model subjects, these approaches tend to require model creators to opt in at least for now. Consent, though enacting a more limited form of participation, *requires* model creators to be wary of unfair and deceptive practices that overstep expressly-informed consent when collecting and using data. Indeed, consent violation is a legally cognizable privacy harm, one with potential repercussions. In the U.S., state information privacy laws do some work to enforce this, with the Illinois Biometric Privacy Act (BIPA) both resulting in a significant number of lawsuits alleging violation, and being responsible for the largest settlement amounts from companies who have breached BIPA by deploying FRT (Strickler, 2020; Yew & Xiang, 2022)). For instance, in *Vance v. IBM*, the court affirmed that IBM violated BIPA by not receiving written consent before collecting and disseminating individuals’ images in their “Diversity in Faces” dataset (Goldenfein, 2023), even though the images used were public. Publicly accessible personal information comes with an intended context of use, which can be violated by memorized and regurgitated data (Carlini et al., 2023). As will be discussed in Section 4.1, to date, the Federal Trade Commission’s enforcement power around unfair and deceptive data practices has centered upon the absence of consent. Practically, this instills consent as the most direct way for data subjects and model subjects to participate in decision-making around the models which affect them, albeit mostly *ex post facto* through their ability to prompt enforcement when discovering that their consent has been violated.

Consent violations frequently occur through improperly scoped consent, where data collected for one purpose is repurposed. This can result in adverse effects beyond the concrete privacy harms (Solove, 2012) that are most often legally enforced in cases such as data breach, e.g., identity theft. As an example, data used beyond its consented purpose leaves data subjects at risk of discrimination harms, facing miscategorization and expansion of surveillance, as detailed above. Such data can also be sold and shared with third parties, further increasing the odds that it is not be-

ing used for the purpose it was collected, and therefore that it is frequently in violation of the consent of data subjects (Calo, 2011). Even when consent is nominally obtained, transparency is often in name only, with data subjects overwhelmed by opaque and all-encompassing digital policies, terms, and conditions (Pasquale, 2021).

Ultimately, the question of consent is complex. Brown et al. (Brown et al., 2022) argue that the current paradigm of training on publicly accessible data makes it highly challenging to distinguish what public data was made public with blanket versus contextual consent, and that, therefore, obtaining informed consent is difficult at best. They make the case for training solely on data explicitly consented for public dissemination. We will return to this argument in Section 5, as it supports using synthetic data generated from properly-consented real data or responsibly procedurally created data, and criticizes using synthetic data generated and used in a manner that exacerbates concerns around consent.

### 2.3. Synthetic Data

Synthetic data in machine learning is defined by its driving goal of mimicking real-world data — it is synthesized to be used as though it were real data for training machine learning algorithms (Jacobsen, 2023). It differs from what is usually referred to as ‘data’, i.e. non-synthetic data, in that it does not have an explicit 1:1 real-world referent. When training a computer vision model to recognize a face, the data traditionally used are representations of real faces, photographs taken of real people. The same holds true for other forms of data — scientific data records representations of physical phenomena such as sensor readings, natural language data is text composed by a real person, etc. Synthetic data is made to resemble these things, but is not explicitly a representation of a real thing. Using the example of a face, a synthetic face could be a drawing that looks for all intents and purposes like a face, but that does not represent a specific person.

In actually creating synthetic data, however, things become muddier. The term encompasses data generated by generative models, more traditionally augmented data, and procedurally created data (Liu et al., 2021). We differentiate between the first two and the latter category based on how derivative of a real-world training dataset they are. *Generated data* is the output of generative models: ML systems that produce a (supposedly novel (Carlini et al., 2023)) output from an input by abstracting over their training data. One generative model that has captured popular attention (Perez, 2023) is StableDiffusion, which generates art from a user-provided input sentence (Cao et al., 2023). *Augmented data* is fuzzier, but equally derivative of a real-world training dataset; the term tends to refer to any real-world data to

which modifications have been made. A model creator seeking to increase the performance of a model on its specific task may create many versions of each image sampled from the input dataset, creating augmented data. This type of synthetic data cannot be considered inherently private or unbiased, with generative models explicitly being found to frequently regurgitate memorized training data (Bai et al., 2021). At larger scales, this type of data can be used as training and evaluation datasets too. As detailed by Khan & Hanna (Khan & Hanna, 2022), datasets are vital components the larger model development cycle, priming synthetic data to reinforce and scale skewed values and requirements that are embedded within models, datasets, and benchmarks.

Generated and augmented data differs from techniques for *procedural creation* of data, where dataset designers make active decisions to create ‘net-new’ representations of data similar to what might be found in the natural world. There are some important differences. The easiest way to visualize how procedural creation works is considering video game character creation — a player starts with a base face shape, and adds the features they want. Notably, with procedural creation of faces, the base volumetric face scans that are used are very far removed from the people that were scanned — they are low fidelity representations, making data lineage even murkier. Instead of outputs being generated from, and therefore bound by training data, procedural creation can result in ‘net-new’ data that never existed previously. Or, as an example outside of computer vision and facial recognition, consider procedurally created finance data which uses agent-based models to mimic real world data generation by creating representative agents and attempting to model money laundering (Lopez-Rojas & Axelsson, 2012). It must be remembered, however, that both agent- and procedural-based synthetic data are highly determined by preconfiguration and design. As such, making inferences about the real world based on procedurally created data is difficult at best. Recalling Jordan et al.’s framework of synthetic’s data usability, this type of synthetic data faces utility and fidelity hurdles (Jordan et al., 2022).

Synthetic data has frequently been explored as a method to avoid privacy concerns, increase model performance and to reduce model bias. Privacy concerns have the longest history of motivating synthetic data (Jordan et al., 2022). Healthcare (Gonzales et al., 2023) and financial (Assefa et al., 2021) domains have been particularly attracted to synthetic solutions due to the sensitivity of their data. Examples including simulation studies in population health (Ngufor et al., 2019); synthetic clinical records used for IT development, education, and training (Davis et al., 2010); money-laundering detection (Lopez-Rojas & Axelsson, 2012); and public release of augmented financial and healthcare datasets to enable open science and research (Harron et al., 2016). In contexts of societal bias, synthetic data has

been explored as a way to remove disparate impact (Feldman et al., 2015; Kamiran & Calders, 2009; Zhang et al., 2016), to suppress imbalance effects and to racially balance datasets (Kortylewski et al., 2019), as well as to remove sensitive information and blind models to race (Wang & Huang, 2019). However, the latter has been found to not always be effective in practice, with applying a ‘veil of ignorance’ not having any notable influence on accuracy of FRT on under-represented categories (Wehrli et al., 2022).

Increasing model performance by using synthetic data has usually meant enlarging datasets to provide robustness to outliers (Wong et al., 2016; Fawaz et al., 2018; Dai et al., 2017). Additionally, operating at a slightly different scale of enlargement (from 0), it has been used in situations where real world data is difficult to access. Google recently demonstrated AlphaGeometry, an AI system purported to solve “Olympiad geometry problems at a level approaching a human-gold medalist”, trained solely on a dataset of 100 million synthetic math proofs — a dataset which could not exist using human generated proofs (Trinh et al., 2024). Computer vision systems require (often impossibly) large amounts of labeled training data in a specific domain (Birhane & Prabhu, 2021).

Finally, consider the context of facial recognition. Though large, open datasets specifically developed for facial recognition tasks exist, such datasets are either extremely basic (i.e. passport photos with great lighting), too narrow (only contain a biased subset of race, gender, head/body pose, etc.), or simply contain glaring and challenging shortcomings (Raji & Fried, 2021). As a result, dataset creators utilize synthetic data. In FRT, this is either (1) procedurally created ‘synthetic’ data, where a bone structure scan is used as a basis for volumetric face models, and then textures representing features are stretched across that model and swapped out (Yi et al., 2014), or (2) the perturbation of existing data to produce more diverse samples from the existing distribution, including both simple techniques and advanced techniques such as diffusion models and generative adversarial networks (GANs) (Dhariwal & Nichol, 2021). We expand upon the use of synthetic data for FRT in the following section.

### 3. Risk 1: Diversity-Washing

This section presents an example of using synthetic data; first describing a dataset that is partially-synthetic, a blend of augmented data and real data, and then describing a dataset that is procedurally created synthetic data. In both cases, we detail how using synthetic data risks creating datasets — and subsequently training and evaluating models from that data — that fail to mitigate bias in data distribution and representation. Furthermore, there is a risk of propagating harm through a patina of legitimacy, and through diversity-

washing potentially harmful models.

The example we use is a real world example, where one of the authors had previously created synthetic datasets to evaluate facial recognition technology (FRT) (Norman et al., 2023). We present this example to illustrate a risk of synthetic data and ground it in a real world setting — this is not an attempt to present novel work on FRT evaluation.

To provide a brief background: FRT is created with the aim of matching images of identifying faces. Companies that sell facial recognition such as Clearview often tout accuracy rates of their systems of 97% or more (Taboh, 2021) — but these calculations are often made under ideal conditions. In real world conditions, such as surveillance camera footage, captured images of faces may be poor quality. FRT has been shown to be prone to make erroneous matches (i.e. identifying someone incorrectly as a match) when using low quality images as input (Hu et al., 2021). However, users of these systems, such as the police, and adjudicators such as judges or members of Congress, who are not experts in FRT or ML/AI, are at a distinct disadvantage in evaluating companies’ claims.

#### 3.1. Augmented Partially-Synthetic Dataset

In order to evaluate FRT models in real world settings, first, benchmark performance for FRT models needed to be established. This occurred by stimulating the (highly unreliable) process of a human identifying an individual from a visual lineup of other humans with similar characteristics. To do so, a source image of a selected identity was identified from the base dataset, detailed below, and a “digital lineup” of (mathematically) similar faces from that base dataset were created. Augmented data was then created by progressively degrading the image of the source identity, and then this augmented data was compared to the similar identities in the digital lineup, as well as to the source, in order to mimic real world settings. The success of the evaluated models was defined by the rate the correct identity was selected with the augmented data, the degraded source images, as input.

A significant body of knowledge already exists concerning both the obvious and non-obvious potential harms in gathering image data containing human subjects, and the real harms of processing such information through FRT (Raji & Fried, 2021). As such, it was important to begin with core datasets that had already been evaluated thoroughly in the literature, rather than collect wholly new human subject data. As such, CASIA-Webface, one of the two most widely used and evaluated public datasets (Yi et al., 2014; Kawulok et al., 2016), was selected as the dataset for use as both non-augmented data and as the base for augmented (in this case, degraded) data. This was chosen due to its sourcing from crawled and scraped publicly available images of celebrities, strict rules prohibiting commercial use, wide variation in



image quality, large number of identities (depth), and large number of images per identity (width) — important features for the dataset that was the provenance for later augmentation. As detailed in Section 2.1, this choice of base dataset is inherently political despite frequently being rendered neutral. Using it as the base for the creation of synthetic data makes it inherently more so due to the downstream effects of using the base dataset in generating derivative data.

The data augmentation techniques used to generate the augmented portion of this mixed dataset can differentially and unpredictably scale issues — making images black and white, as an example, could further segment training data by racial presentation. Model creators frequently augment data while training models, past the stages of development where they are considering dataset collection and annotation. This process has a set of well documented risks for model fairness. However, in using this process to create large base datasets, there is a change of framing that re-introduces these risks. For example, the dataset created for the FRT evaluation (created by augmenting data and combining with real data) was created with the explicit goal of being more representative of real world conditions. Datasets are frequently treated as ground truth (Birhane & Prabhu, 2021), hiding the decisions and processes by which they were created. This risks ignoring issues that can occur from augmenting data. Even if synthetic data appears 'diverse', the generation of that data cannot be unwound from the particular datasets and models that it is being generated from, and any attendant shortcomings. To start, any biased representations would at best replicate from the original dataset. If data augmentation technique(s) impacted some subjects differently than others, the resulting impact could be unintended bias in the dataset. Since the presence of such relationships are rarely known, much less understood statistically in datasets, it is also possible that the sampling strategy used to choose which data from the dataset will be used in training may actually exacerbate harm by over-representing biased representations. Deep learning techniques are generally already susceptible to overfitting, where a model learns how to predict patterns in a way that pays too much attention to the training data, and doesn't generalize to other data — it learns specific idiosyncrasies and meaningless data artifacts. Synthetic data seems like it should have the capacity to remedy overfitting, through careful and bespoke dataset construction that debiases data distributions. One could assume this would increase fidelity and enable better generalization over a more diverse training space. In reality however, when synthetic data is overfitted, these idiosyncrasies can go through the entire model training process unnoticed. As such, synthetic data instead increases the likelihood of overfitting errors being propagated through, necessitating that further technical care is taken to prevent overfitting. In our FRT evaluation example, such preventative measures were taken by curating

the dataset so that visible artifacts such as skin tone were at parity with acceptable real world dataset distributions. However, the perils of overfitting are a way in which synthetic data can struggle to meet the standard of utility necessary to work as a replacement for real world data (Jordon et al., 2022).

### 3.2. Procedurally Created Fully Synthetic Dataset

Beyond the above partially-synthetic dataset, there was a need to better evaluate performance on specific types of data not present in the original datasets that the FRT models were trained on. So, a synthetic dataset consisting of procedurally created data, namely mixed examples of non-degraded and degraded computer-generated faces (Qiu et al., 2021; Basak et al., 2021), was developed. As previously detailed, this is a common usage of synthetic data — needed representative data was not available for collection, and so generating synthetic data was the easiest method of proceeding.

The Synthesis.AI software<sup>2</sup> used to create the fully synthetic dataset (as with most procedural synthetic human/object generation tools) works by providing unprecedented control over how a dataset and its inherent metadata parameters are specified. This software employs a combination of classic rendering and generative synthesis to create photo-realistic images of human faces, bodies, and environments (Nikolenko, 2022). A user is able to decide how much and which type of each characteristic (in our case age, race, gender, hair type, pose, lighting etc.). However, the tool did not make any suggestions or place any controls based on sociotechnical norms or demographic data (such as the census etc.) when creating a synthetic human dataset of any type. When first testing the Synthesis.AI API, a dramatically racially imbalanced dataset was returned, even though the specification given was for randomization of the race characteristic. At first glance, the dataset appeared diverse and was numerically at parity for gender. However, the software lacked permutations for Asian people, Middle Eastern people and Black women, leading to a stark racial disparity upon deeper inspection, and a preponderance of white men and white women despite attempts at balancing racial demographics. Such a system allows any user to easily create an unintentionally biased dataset, which could then be used to train a biased model. Instead of mitigating data distribution and representation concerns, this risks extending them.

As a further example, Microsoft's FaceSynthetics (Wood et al., 2021) is a procedurally created synthetic dataset of 100,000 individuals, with faces derived from representative 511 base scans. However, these 511 base scans include only 30 Black men, and even fewer Hispanic/Arab/Indian men and Black/Hispanic/Arab/Indian women (borrowing the reported demographic categories), meaning that the fully

<sup>2</sup>Synthesis AI (<https://synthesis.ai>)

diverse population they claim include multiple racial categories fully defined by the ways in which these 30 faces can be manipulated through a generative process. These manipulations include fine-tuning hair, expression, and clothing, but published details on the process of how these potentially racially-coded aspects were chosen are sparse. It is not known how those features are distributed in real faces, and attempting to extrapolate a portion of a representatively diverse dataset from such a small set of base faces leads to a risk of *statistical diversity without representational diversity*, compared to a representative dataset of real images with both statistical and representational diversity. Synthetic data here falls flat in addressing these complex, cultural and deeply contextualized factors.

These tools risk falling into the ‘panacea of legitimization’ that Frank Pasquale describes (Pasquale, 2021), where ethical concerns are not only routed around, but are routed around in such a manner that they can reify malpractice due to the co-constitutive nature of ML practices and computing platforms (Berman, 2023). We point to recent work focused on toolkits for supporting practitioners in contextualizing ML system work as an avenue for improving upon this (Deng et al., 2022).

#### 4. Risk 2: Circumvented Consent

Consent has become a key component of privacy and data protection, both via regulatory enforcement and as a necessary foundation for the participatory approaches that have emerged as ethical practice for preventing harm. Consent also plays a role in U.S. sectoral statutes such as HIPPA, U.S. state laws such as California’s California Privacy Rights Act and Illinois’ Biometric Information Privacy Act, as well as laws in many other countries, with the EU’s GDPR a notable example. These statutes share a common goal: to provide people with control over their personal data, via notification, access, and consent regarding the collection, use, and disclosure of personal data. This type of privacy regulation is referred to as “privacy self-management” by Solove (Solove, 2012), and focuses solely on whether or not data subjects have consented, rather than on value judgements of privacy practices. This section will focus on illustrating the risk that synthetic data poses to consent-based frameworks by expanding upon how the Federal Trade Commission (FTC) has functionally used consent as a key aspect of conducting enforcement against companies using ML systems. The analysis is guided by one author’s experience at the FTC, but draws upon solely public knowledge.

#### 4.1. Consent, Deception and Model Deletion

The FTC plays a vital role in the current U.S. privacy legal framework. This framework emphasizes individuals’ notice of, and consent to, the collection and use of their data. The FTC is an independent agency of the United States government that is tasked with protecting consumers and promoting competition in the marketplace. In the absence of federal privacy law, the FTC has played the role of de facto privacy enforcement, primarily based on its authority to police unfair and deceptive business practices. The Federal Trade Commission Act, and specifically Section 5, is a broadly applicable federal statute prohibiting “unfair or deceptive acts and practices”<sup>3</sup>. An *unfair* practice “causes or is likely to cause substantial injury to consumers which is not reasonably avoidable by consumers themselves and not outweighed by countervailing benefits to consumers or to competition”, while a *deceptive* practice includes “any ‘representation, omission, or practice’ that is (i) material, and (ii) likely to mislead consumers who are acting reasonably under the circumstances” (Solove & Hartzog, 2014). Notably, deception does not require any proof of intent. The FTC has brought deception claims against companies who have violated the terms of their privacy policies, failed to uphold promises of data security, or have failed to provide sufficient notice regarding data collection and use (Solove & Hartzog, 2014).

In settling cases against companies that have deceptively collected data, the FTC has required not only that the data in question be deleted and the affected users be notified, but also that all “affected work product” (Li, 2022) be deleted as well — including models trained on that data. This approach is referred to as *model deletion*. The FTC posits that this approach is necessary in order to ensure that companies do not profit from the unfair or deceptive collection of data, and to prevent them from using the data in the future. Intellectual property is typically a tech company’s most valuable asset; it is an important factor for securing venture capital funding and the sale or licensing of IP often comprises tech companies’ core business models. In forcing a company to delete models, the FTC has also significantly changed the deterrence calculus for companies: from paying (relatively) small fines, to potentially losing a vital business asset (Elder, 2022). The FTC has used the concept of model deletion in recent enforcement actions against companies that have collected data deceptively, including Amazon Ring, RiteAid, Everalbum, Clearview and Kurbo (WeightWatchers) (Hutson & Winters, 2022).

<sup>3</sup>15 U.S.C. Sec. 45(a)



## 4.2. Synthetic Data and FTC Enforcement

The use of synthetic data risks undermining the utility of deception-based enforcement in regulating the collection of data, and therefore also undermines the regulation of models trained on such synthetic data. As previously described in Section 2.1, datasets play a foundational role in the models trained on them, and trusting models trained on deceptively collected datasets to operate without harm seems foolhardy. Enforcement by the FTC has hinged upon arguments that data was collected and used deceptively — often argued due to the absence of proper consent. By using synthetic data, however, it becomes easy for model creators to obfuscate the origins and consent of the data being used to create models. In the case of a procedurally created synthetic dataset, consent is no longer a procedural hook to limit downstream harms flowing from use, while in other synthetic datasets, unless data lineage is carefully recorded, traceability to the original data is at risk (Scheuerman et al., 2023).

Synthetic data also exacerbates existing logistical challenges for model deletion as an enforcement tool. Synthetic data brings questions of data lineage to the forefront, as ever-more-complicated sets of original, augmented, and derivative data are produced based on new face datasets with millions of people. As a small example, keeping track of whether a single version of a dataset has undergone ethical testing, or was sectioned off as a test dataset, is a challenge for FTC enforcers — let alone when datasets include different scaling factors and different degradations, with different subsets of identities (generated, procedurally created or real) and different levels of augmentation. One of the key hurdles to model deletion is the requirement for a high level of internal company documentation and logging. This documentation and logging is essential to identify the data that was collected illegally or deceptively, as well as the work product that was developed using that data. However, not all companies have robust internal documentation and logging systems, which can make it difficult for the FTC to determine the extent of the harm caused by the illegal or deceptive data collection practices. Another challenge is authentication and audit. Companies must demonstrate that they have successfully deleted the affected work product, and the FTC or other enforcement bodies must have a way of verifying that the company is being honest. However, this can be difficult, as it requires a level of trust in the company and its processes. In a setting where the FTC has to this point relied on settlement agreements, synthetic data further complicates existing logistical challenges, presenting important friction to enforcement.

## 4.3. Beyond Deception

In considering enforcement against companies using ML systems, it is important to note that beyond deception, the FTC has also enforced its unfairness authority. This occurred in the case of RiteAid, where a biased facial recognition model was used to falsely identify people in certain protected classes as more likely to commit crime. This case was first-in-kind, but demonstrates that the FTC is not solely beholden to deceptive data collection as an avenue for enforcement. However, the details of the case were particularly egregious, with RiteAid failing to undertake even the most basic risk assessments, and in part hinged on a violation of a previous settlement. Additionally, the system was trained on in-store camera footage without consent from data subjects, and model subjects were not notified or able to opt-out. Thus if RiteAid had been investigated for deceptive data collection, harm resulting from this system could have been prevented at the point where the system was trained non-consensually. But what would happen if RiteAid had trained its FRT model on synthetic data? FTC enforcement would have to hinge on unfair practice alone. While successful here, the case against RiteAid was egregious. The Supreme Court’s neutering of the FTC’s power to levy fines (Chopra & Levine, 2021/2022), in addition to both the deception and unfairness enforcements occurring through settlement rather than being decided in court, means the boundaries of the FTC’s ability to engage in this type of enforcement are still undefined. As such, the FTC’s ability to intervene both at the data collection stage (deception) and the model deployment stage (unfairness) gives options<sup>4</sup>. Synthetic data complicates the usage of a demonstrably useful tool for protecting data subjects and model subjects, by complicating the use of the deception standard.

Finally, it cannot be forgotten that while risking increased friction and obfuscation, synthetic datasets composed of augmented or generated data are demonstrably derivative, inherently based on real data representing real data subjects. And, despite the veneer provided by language such as ‘net-new’, procedural creation of synthetic data is also derivative, and thus suffers from issues of consent and participation too. In the example discussed in Section 3.2, the procedurally created synthetic dataset for FRT evaluation, the dataset was generated using Synthesis.AI’s commercial software. Software tools such as Synthesis.AI often utilize face and body scanning technology as the foundation of their generative processes, raising concerns around the limits of informed consent. The data subjects upon which these technologies are trained are rendered invisible and thus the use of such software is predicated on, at best (Raji

<sup>4</sup>While there exists the potential for both unfairness at data collection and deception at model deployment, cases to-date have lined up in this fashion

et al., 2020), ambiguous consent. Similarly, it is important to acknowledge that the data subjects whose likenesses are captured in the CASIA-Webface, while primarily scraped from public sources, were not asked for informed consent regarding their data’s use — even when it has been shown that having your face included in such a dataset increases the accuracy of facial recognition models on your specific face (Dulhanty & Wong, 2020). As detailed by Peng et al. (Peng et al., 2021), using derivatives of common datasets introduces scaling concerns around propagation of improperly consented data, and as such, using synthetic data risks *further* scaling propagation of this issue. In decoupling data subjects from their data, this also removes their capacity to participate. Reconsider the Participatory Data Stewardship Framework mentioned above in Section 2.2; all five levels require that data subjects have at the very least visibility, and preferably control, over their data (Patel et al., 2021). In further removing the ability for data subjects to consent, not only is that minimal level of agency reduced, but the potential for involvement in decision-making that directly effects them is erased.

## 5. Discussion

The positioning of synthetic data as a panacea to problems of representation and deceptive data collection, furthered by its portrayal as synthetic, as neutral, as created without lineage, risks placing the means of fixing those problems in the hands of those who created them and trusting that they’ll get it right. Instead, the kinds of racialized misrecognition and bias that Ruha Benjamin, Safiya Noble, and others have drawn attention to must be considered when determining whether to use synthetic data. As Benjamin argues, our current machine learning development ecosystem must reckon with a history of discriminatory design in which racist values and assumptions are built into our technical systems. The ‘new Jim Code’, as she terms it, works to deepen the production of disparate harm, even while cloaked in neutrality and the language of innovation (Benjamin, 2019; Irani, 2019). Discriminatory practices are inherent to the current state of AI system development, privileging whiteness and discriminating against people of color, specifically women of color (Noble, 2021). Sara Ahmed’s work on the phenomenology of whiteness highlights the danger of a solution that further enables a reification of non-whiteness as a space outside. She details that “institutional spaces are shaped by the proximity of some bodies and not others: white bodies gather, and cohere to form the edges of such spaces” (Ahmed, 2007). Synthetic data as a fix in this racialized context risks further enabling amplification of racial hierarchies, allowing for those within the boundaries to actively constitute the exclusionary and weaponized edges of these spaces: to define a face, train a model based on that definition, and decide its performance based on label-

ing racial boundaries. It risks not alleviating but instead contributing to *race as a technology*, designed to “stratify and sanctify” social injustice in the architecture of everyday life (Benjamin, 2019); an added consolidation of power.

Another example of the risks of consolidation of power through synthetic data arises when considering the inherent relational aspects of data privacy. Solon Barocas and Helen Nissenbaum identify the risk of a “tyranny of the minority” in big-data analytics when “the volunteered information of the few can unlock the same information about the many” (Barocas & Nissenbaum, 2014). More recently, Salome Viljoen emphasizes the importance of a relational theory of data governance (Viljoen, 2021). As Viljoen explains, dataflows entail not only ‘vertical’ relations between a particular individual and a data collector, but also ‘horizontal’ relations between the individual and others sharing relevant population features. Viljoen focuses on the manner in which informational infrastructures rely on group classification to make sense of individuals by taking a ‘relevant shared feature,’ generating a prediction based upon that shared feature, and then applying this prediction. When those shared features are derived from synthetic data, decoupled from any real context and perhaps even specifically created to rectify gaps in representation, we hand power to those creating that synthetic data. We risk imposing designers’ decision-making in lieu of and upon those least likely to have been represented and most likely to be harmed by both the diversity-washing and the side-stepping of consent. After all, if they were represented or able to consent in the first place, there would be no need for additional synthetic data. Data minimization and lineage principles (Hutson & Winters, 2022) are a first step towards mitigating this issue by requiring documentation and its requisite transparency into where data has come from. The need for this is also readily apparent when considering contexts such as the EU’s Right to be Forgotten (Razmetaeva, 2020), where synthetic data further complicates the ability to be removed from a dataset. In making it harder to decouple data from its context through the use of synthetic data, there is an avenue for mitigating consolidation of power and ensuring consent.

Many of the risks discussed in this paper propagate from the centralized decision-making nature that synthetic data imposes. Participatory governance structures, as mentioned in Section 2, offer a potential solution here — synthetic data could be created to represent concerned groups *by* those self-same groups, re-establishing control and mitigating some concerns around consent and contextualization. Many have called for training data to be restricted to only data that is explicitly consented to be used, though consent is difficult (if not impossible) to establish and propagate over multiple degrees of separation (Brown et al., 2022). But synthetic data, when generated to purpose by concerned communities, can provide a potential solution. Models such as those pre-

sented in the field of Indigenous Data Sovereignty, where there has been effective push back against external categorization schemas (Rainie et al., 2019) show potential for participatory governance models to address group misrepresentation (Andrus & Villeneuve, 2022).

Additionally, there are practical considerations that make ‘participatory synthetic data’ an attractive path forward. Both language and computer vision models are beginning to contend with a phenomena commonly referred to as ‘garbage-in garbage-out’ (Shumailov et al., 2023; Martínez et al., 2023). This refers to the advent of generated data becoming commonplace and public, and the related struggles by those capturing data to differentiate between that data and real data, leading to data generated by a system such as ChatGPT becoming its own training data in the future. Work by Agnew et al. (Agnew et al., 2024) examines the use of these models to replace participants in industry research, highlighting how proposals to do so are motivated by cost reduction and data diversity. They identify these proposals as facing issues in aligning with the values human participants identify as important, specifically including inclusion and representation, necessitating further contextualization and bespoke dataset creation. In such a world, large tech companies may have business motivations for engaging with responsibly created synthetic data, and as demonstrated by Deng et al. (Deng et al., 2022), methods exist for enabling machine learning practitioners to better contextualize the work they do — a vital aspect of any future responsible synthetic data work.

Further responsible dataset development frameworks that explicitly attend to the particulars of synthetic data, as well as tooling and practice that examines and makes transparent the provenance of synthetic data, are needed. As examples of this, we propose ‘how could less risky synthetic data be produced?’, as well as ‘how could governance approach consent issues with synthetic data?’ as important future research questions. In future work, we hope to follow the call of Denton et al. (Denton et al., 2021), contesting machine learning datasets and focusing on the “contingent, historical, and value-laden work practices of actual machine learning researchers” to better understand how the practice of using synthetic data is motivated, the contingent conditions that have lead to its common usage, and the norms and routines that surround it. In so doing, there is the opportunity to survey and better understand the use of synthetic data and create better tools and frameworks for both mitigating its potential for harmful power consolidation, as well as to envision how it can be used as a tool for taking power back (Whitney et al., 2021).

## 6. Conclusion

In this paper, we build on prior responsible dataset development work by focusing on the under-explored impacts of synthetic data on dataset development. Synthetic data will continue to play an ever-increasing role in the training of machine learning systems as real-world data becomes harder to capture, and we must attend to language that paints it as a panacea. We show two examples of the risks of synthetic data, diversity-washing and consent circumvention, and discuss how it is a complicated tool, gravitationally prone to consolidation of power, but with potential for being used to enable participatory governance instead of squashing it.

## 7. Researcher Ethics and Social Impact

### 7.1. Researcher Positionality Statement

The first author is a white Latino AI researcher significantly influenced by their research, which has examined how policy and technical practice interplay and talk past each other, and how this dynamic affects those most likely to be harmed by AI systems. They worked in ML before moving into academia. The second author is a Black AI researcher with a variety of experiences in government, industry and now academia. They have access to the resources necessary to conduct their research, and recognize that that they have access to resources that many others do not. They strive to be conscious of their biases and to mitigate their impact on their work as much as possible. Both authors were motivated to write this paper by the realization that the risks of a commonplace technical practice were under-explored when discussing a real world example (detailed in this paper), and hoped to provide a starting point for understanding how using synthetic data could go wrong. Both researchers are based in the U.S., and that heavily influences both the examples they draw upon to show risks, the harms that they find salient, and the Overton window through which they view the world.

### 7.2. Ethical Considerations Statement

This work focuses on illustrating risks through the analysis and description of public-facing information and prior work through a new lens. As this is an example of RAI work that is focused on human impact but that does not involve study participants or create or deploy new technology, the main ethical consideration is in how this prior work is presented, where we actively attempted to avoid falling into some of the same traps we discuss — we do not wish to make the technology seem inevitable or help to legitimize it, while we also do not want to forestall the opportunity for the risks we present to be mitigated and it to be used in participatory and ethical manners. We believe that the FRT evaluation example provided in this paper, created to assist in preventing



unfounded FRT claims being used in the criminal justice setting, necessitated the creation of these datasets and the usage of synthetic data, but that is far from an ever-present conclusion.

### 7.3. Adverse Impact Statement

The largest adverse impact we are wary of is that these risks could be taken as playbooks — we hope that nobody comes away thinking that there is opportunity to take advantage of them. We think that by making them public we are doing more of a good, as these examples demonstrate that the potential for these things already exists, and active exploration and research focused on mitigating and re-directing potential is the best way forward. We also see that this work could potentially draw scrutiny to legitimate uses of synthetic data, but hope that any added friction there is worth preventing potential malpractice.

### 7.4. Acknowledgements

This research was supported by UC Berkeley’s Center for Long Term Cybersecurity AI Policy Hub, where the first author conducted research as a Fellow. This research was also supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE 1752814. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or of the CLTC. We thank those who helped us workshop and develop this material, specifically including Professor Deirdre Mulligan, Professor Jenna Burrell, Professor Hany Farid, Professor Brandie Nonnecke, Jessica Newman, Lauren Chambers and Seyi Olojo.

## References

- Abebe, R. and Goldner, K. Mechanism design for social good. *AI Matters*, 4(3):27–34, October 2018. ISSN 2372-3483. doi: 10.1145/3284751.3284761.
- Agnew, W., Bergman, A. S., Chien, J., Díaz, M., El-Sayed, S., Pittman, J., Mohamed, S., and McKee, K. R. The illusion of artificial inclusion, February 2024.
- Ahmed, S. A phenomenology of whiteness. *Feminist Theory*, 8(2):149–168, August 2007. ISSN 1464-7001, 1741-2773. doi: 10.1177/1464700107078139.
- Alkhatib, A. and Bernstein, M. Street-Level Algorithms: A Theory at the Gaps Between Policy and Decisions. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, Glasgow Scotland Uk, May 2019. ACM. ISBN 978-1-4503-5970-2. doi: 10.1145/3290605.3300760.
- Andrus, M. and Villeneuve, S. Demographic-Reliant Algorithmic Fairness: Characterizing the Risks of Demographic Data Collection in the Pursuit of Fairness. In *2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT ’22*, pp. 1709–1721, New York, NY, USA, June 2022. Association for Computing Machinery. ISBN 978-1-4503-9352-2. doi: 10.1145/3531146.3533226.
- Arnstein, S. R. A Ladder Of Citizen Participation. *Journal of the American Institute of Planners*, 35(4):216–224, July 1969. ISSN 0002-8991. doi: 10.1080/01944366908977225.
- Assefa, S. A., Dervovic, D., Mahfouz, M., Tillman, R. E., Reddy, P., and Veloso, M. Generating synthetic data in finance: Opportunities, challenges and pitfalls. In *Proceedings of the First ACM International Conference on AI in Finance, ICAIF ’20*, pp. 1–8, New York, NY, USA, October 2021. Association for Computing Machinery. ISBN 978-1-4503-7584-9. doi: 10.1145/3383455.3422554.
- Bai, C.-Y., Lin, H.-T., Raffel, C., and Kan, W. C.-w. On Training Sample Memorization: Lessons from Benchmarking Generative Modeling with a Large-scale Competition. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, KDD ’21*, pp. 2534–2542, New York, NY, USA, August 2021. Association for Computing Machinery. ISBN 978-1-4503-8332-5. doi: 10.1145/3447548.3467198.
- Barnett, J. and Diakopoulos, N. Crowdsourcing Impacts: Exploring the Utility of Crowds for Anticipating Societal Impacts of Algorithmic Decision Making. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 56–67, July 2022. doi: 10.1145/3514094.3534145.
- Barocas, S. and Nissenbaum, H. Big data’s end run around procedural privacy protections. *Communications of the ACM*, 57(11):31–33, October 2014. ISSN 0001-0782. doi: 10.1145/2668897.
- Basak, S., Corcoran, P., Khan, F., McDonnell, R., and Schukat, M. Learning 3D Head Pose From Synthetic Data: A Semi-Supervised Approach. *IEEE Access*, 9: 37557–37573, 2021. ISSN 2169-3536. doi: 10.1109/ACCESS.2021.3063884.
- Benjamin, R. *Race after Technology: Abolitionist Tools for the New Jim Code*. Polity, Medford, MA, 2019. ISBN 978-1-5095-2643-7.
- Berman, G. Machine Learning practices and infrastructures. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society, AIES ’23*, pp. 466–481, New York, NY, USA, August 2023. Association

- for Computing Machinery. ISBN 9798400702310. doi: 10.1145/3600211.3604689.
- Birhane, A. and Prabhu, V. U. Large image datasets: A pyrrhic win for computer vision? In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1536–1546. IEEE Computer Society, January 2021. ISBN 978-1-66540-477-8. doi: 10.1109/WACV48630.2021.00158.
- Birhane, A., Isaac, W., Prabhakaran, V., Díaz, M., El-ish, M. C., Gabriel, I., and Mohamed, S. Power to the People? Opportunities and Challenges for Participatory AI. In *Equity and Access in Algorithms, Mechanisms, and Optimization*, pp. 1–8, October 2022. doi: 10.1145/3551624.3555290.
- Bivens, R. The gender binary will not be deprogrammed: Ten years of coding gender on Facebook. *New Media & Society*, 19(6):880–898, June 2017. ISSN 1461-4448. doi: 10.1177/1461444815621527.
- Bowman, S. R. and Dahl, G. E. What Will it Take to Fix Benchmarking in Natural Language Understanding?, October 2021.
- Bratteteig, T. and Wagner, I. Unpacking the Notion of Participation in Participatory Design. *Computer Supported Cooperative Work*, 25(6):425–475, December 2016. ISSN 0925-9724. doi: 10.1007/s10606-016-9259-4.
- Brown, H., Lee, K., Mireshghallah, F., Shokri, R., and Tramèr, F. What Does it Mean for a Language Model to Preserve Privacy? In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pp. 2280–2292, Seoul Republic of Korea, June 2022. ACM. ISBN 978-1-4503-9352-2. doi: 10.1145/3531146.3534642.
- Bullock, J. B. Artificial Intelligence, Discretion, and Bureaucracy. *The American Review of Public Administration*, 49(7):751–761, October 2019. ISSN 0275-0740. doi: 10.1177/0275074019856123.
- Buolamwini, J. and Gebru, T. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, pp. 77–91. PMLR, January 2018.
- Calo, R. The Boundaries of Privacy Harm Essay. *Indiana Law Journal*, 86(3):1131–1162, 2011.
- Cao, H., Tan, C., Gao, Z., Xu, Y., Chen, G., Heng, P.-A., and Li, S. Z. A Survey on Generative Diffusion Model, December 2023.
- Carlini, N., Hayes, J., Nasr, M., Jagielski, M., Sehwag, V., Tramèr, F., Balle, B., Ippolito, D., and Wallace, E. Extracting Training Data from Diffusion Models, January 2023.
- Chamikara, M. A. P., Bertok, P., Khalil, I., Liu, D., and Camtepe, S. Privacy Preserving Face Recognition Utilizing Differential Privacy. *Computers & Security*, 97: 101951, October 2020. ISSN 0167-4048. doi: 10.1016/j.cose.2020.101951.
- Chopra, R. and Levine, S. A. The Case for Resurrecting the FTC Act’s Penalty Offense Authority. *University of Pennsylvania Law Review*, 170(1):71–124, 2021/2022.
- Christiano, P., Leike, J., Brown, T. B., Martic, M., Legg, S., and Amodei, D. Deep reinforcement learning from human preferences, February 2023.
- Christin, A. Algorithms in practice: Comparing web journalism and criminal justice. *Big Data & Society*, 4(2): 2053951717718855, December 2017. ISSN 2053-9517. doi: 10.1177/2053951717718855.
- Corbett, E., Denton, E., and Erete, S. Power and Public Participation in AI. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO ’23, pp. 1–13, New York, NY, USA, October 2023. Association for Computing Machinery. ISBN 9798400703812. doi: 10.1145/3617694.3623228.
- Crawford, K. and Paglen, T. Excavating AI: The politics of images in machine learning training sets. *AI & SOCIETY*, 36(4):1105–1116, December 2021. ISSN 1435-5655. doi: 10.1007/s00146-021-01162-8.
- Dai, Z., Yang, Z., Yang, F., Cohen, W. W., and Salakhutdinov, R. Good Semi-supervised Learning that Requires a Bad GAN, November 2017.
- Davis, P., Lay-Yee, R., and Pearson, J. Using micro-simulation to create a synthesised data set and test policy options: The case of health service effects under demographic ageing. *Health Policy (Amsterdam, Netherlands)*, 97(2-3):267–274, October 2010. ISSN 1872-6054. doi: 10.1016/j.healthpol.2010.05.014.
- Deng, W. H., Nagireddy, M., Lee, M. S. A., Singh, J., Wu, Z. S., Holstein, K., and Zhu, H. Exploring How Machine Learning Practitioners (Try To) Use Fairness Toolkits. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pp. 473–484, Seoul Republic of Korea, June 2022. ACM. ISBN 978-1-4503-9352-2. doi: 10.1145/3531146.3533113.
- Denton, E., Hanna, A., Amironesei, R., Smart, A., and Nicole, H. On the genealogy of machine learning datasets: A critical history of ImageNet. *Big Data & Society*, 8(2): 20539517211035955, July 2021. ISSN 2053-9517. doi: 10.1177/20539517211035955.

- Dhariwal, P. and Nichol, A. Diffusion Models Beat GANs on Image Synthesis, June 2021.
- Diaz, M., Kivlichan, I. D., Rosen, R., Baker, D. K., Amironei, R., Prabhakaran, V., and Denton, E. Crowd-WorkSheets: Accounting for Individual and Collective Identities Underlying Crowdsourced Dataset Annotation. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pp. 2342–2351, June 2022. doi: 10.1145/3531146.3534647.
- Dulhanty, C. and Wong, A. Investigating the Impact of Inclusion in Face Recognition Training Data on Individual Face Identification, January 2020.
- Elder, E. Wrongful Improvers as a Guiding Principle for Application of the FTC’s IP Deletion Requirement. *WASHINGTON LAW REVIEW*, 97, 2022.
- Fawaz, H. I., Forestier, G., Weber, J., Idoumghar, L., and Muller, P.-A. Data augmentation using synthetic data for time series classification with deep residual networks, August 2018.
- Feldman, M., Friedler, S., Moeller, J., Scheidegger, C., and Venkatasubramanian, S. Certifying and removing disparate impact, July 2015.
- Gilman, M. E. Beyond Window Dressing: Public Participation for Marginalized Communities in the Datafied Society, November 2022.
- Glucker, A. N., Driessen, P. P. J., Kolhoff, A., and Runhaar, H. A. C. Public participation in environmental impact assessment: Why, who and how? *Environmental Impact Assessment Review*, 43:104–111, November 2013. ISSN 0195-9255. doi: 10.1016/j.eiar.2013.06.003.
- Goldenfein, J. Privacy’s Loose Grip on Facial Recognition - Law and the Operational Image, December 2023.
- Gonzales, A., Guruswamy, G., and Smith, S. R. Synthetic data in health care: A narrative review. *PLOS Digital Health*, 2(1):e0000082, January 2023. ISSN 2767-3170. doi: 10.1371/journal.pdig.0000082.
- Gray, M. L. and Suri, S. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Eamon Dolan Books, 2019.
- Groves, L., Peppin, A., Strait, A., and Brennan, J. Going public: The role of public participation approaches in commercial AI labs. In *2023 ACM Conference on Fairness, Accountability, and Transparency*, pp. 1162–1173, Chicago IL USA, June 2023. ACM. ISBN 9798400701924. doi: 10.1145/3593013.3594071.
- Halloran, J., Hornecker, E., Stringer, M., Harris, E., and Fitzpatrick, G. The value of values: Resourcing co-design of ubiquitous computing. *CoDesign*, 5(4):245–273, December 2009. ISSN 1571-0882. doi: 10.1080/15710880902920960.
- Hamidi, F., Scheuerman, M. K., and Branham, S. M. Gender Recognition or Gender Reductionism? The Social Implications of Embedded Gender Recognition Systems. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI ’18, pp. 1–13, New York, NY, USA, April 2018. Association for Computing Machinery. ISBN 978-1-4503-5620-6. doi: 10.1145/3173574.3173582.
- Harron, K., Gilbert, R., Cromwell, D., and van der Meulen, J. Linking Data for Mothers and Babies in De-Identified Electronic Health Data. *PLOS ONE*, 11(10):e0164667, October 2016. ISSN 1932-6203. doi: 10.1371/journal.pone.0164667.
- Harvey, A. and LaPlace, J. Researchers Gone Wild: Origins and Endpoints of Image Training Datasets Created “In the Wild”. In Herlo, B., Irrgang, D., Joost, G., and Unteidig, A. (eds.), *Design*, volume 54, pp. 289–310. transcript Verlag, Bielefeld, Germany, 1 edition, December 2021. ISBN 978-3-8376-5760-9 978-3-8394-5760-3. doi: 10.14361/9783839457603-016.
- Hu, C., Sapkota, B. B., Thomasson, J. A., and Bagavathiannan, M. V. Influence of Image Quality and Light Consistency on the Performance of Convolutional Neural Networks for Weed Mapping. *Remote Sensing*, 13(11):2140, January 2021. ISSN 2072-4292. doi: 10.3390/rs13112140.
- Hügel, S. and Davies, A. R. Public participation, engagement, and climate change adaptation: A review of the research literature. *Wiley Interdisciplinary Reviews. Climate Change*, 11(4):e645, 2020. ISSN 1757-7780. doi: 10.1002/wcc.645.
- Hutchinson, B., Smart, A., Hanna, A., Denton, E., Greer, C., Kjartansson, O., Barnes, P., and Mitchell, M. Towards Accountability for Machine Learning Datasets: Practices from Software Engineering and Infrastructure. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAcCT ’21, pp. 560–575, New York, NY, USA, March 2021. Association for Computing Machinery. ISBN 978-1-4503-8309-7. doi: 10.1145/3442188.3445918.
- Hutson, J. and Winters, B. America’s Next ‘Stop Model!’: Model Disgorgement, September 2022.
- Irani, L. *Chasing Innovation: Making Entrepreneurial Citizens in Modern India*. Princeton Studies in Culture



- and Technology. Princeton University Press, Princeton, New Jersey ; Oxford, United Kingdom, 2019. ISBN 978-0-691-17514-0 978-0-691-17513-3.
- Irgens, G. A., Adisa, I., Bailey, C., and Quesada, H. V. Designing with and for Youth: A Participatory Design Research Approach for Critical Machine Learning Education. *Educational Technology & Society*, 25(4):126–141, 2022. ISSN 1176-3647.
- Iversen, O. S., Halskov, K., and Leong, T. W. Values-led participatory design. *CoDesign*, 8(2-3):87–103, June 2012. ISSN 1571-0882. doi: 10.1080/15710882.2012.672575.
- Jackson, S. J. Rethinking Repair. In Gillespie, T., Boczkowski, P. J., and Foot, K. A. (eds.), *Media Technologies*, pp. 221–240. The MIT Press, February 2014. ISBN 978-0-262-52537-4. doi: 10.7551/mitpress/9780262525374.003.0011.
- Jacobsen, B. N. Machine learning and the politics of synthetic data. *Big Data & Society*, 10(1): 20539517221145372, January 2023. ISSN 2053-9517. doi: 10.1177/20539517221145372.
- Jordon, J., Szpruch, L., Houssiau, F., Bottarelli, M., Cherubin, G., Maple, C., Cohen, S. N., and Weller, A. Synthetic Data – what, why and how?, May 2022.
- Kamiran, F. and Calders, T. Classifying without discriminating. In *Control and Communication 2009 2nd International Conference on Computer*, pp. 1–6, February 2009. doi: 10.1109/IC4.2009.4909197.
- Kawulok, M., Celebi, M. E., and Smolka, B. (eds.). *Advances in Face Detection and Facial Image Analysis*. Springer International Publishing, Cham, 2016. ISBN 978-3-319-25956-7 978-3-319-25958-1. doi: 10.1007/978-3-319-25958-1.
- Khan, M. and Hanna, A. The Subjects and Stages of AI Dataset Development: A Framework for Dataset Accountability, September 2022.
- Kortylewski, A., Egger, B., Schneider, A., Gerig, T., Morel-Forster, A., and Vetter, T. Analyzing and Reducing the Damage of Dataset Bias to Face Recognition With Synthetic Data. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 2261–2268, Long Beach, CA, USA, June 2019. IEEE. ISBN 978-1-72812-506-0. doi: 10.1109/CVPRW.2019.00279.
- Lee, N. T., Resnick, P., and Barton, G. Algorithmic bias detection and mitigation: Best practices and policies to reduce consumer harms, May 2019.
- Li, T. C. Algorithmic Destruction, 2022.
- Liu, J., Snodgrass, S., Khalifa, A., Risi, S., Yannakakis, G. N., and Togelius, J. Deep learning for procedural content generation. *Neural Computing and Applications*, 33(1):19–37, January 2021. ISSN 1433-3058. doi: 10.1007/s00521-020-05383-8.
- Lopez-Rojas, E. A. and Axelsson, S. Money Laundering Detection using Synthetic Data. In *Annual Workshop of the Swedish Artificial Intelligence Society (SAIS)*. Linköping University Electronic Press, Linköpings universitet, 2012.
- Martínez, G., Watson, L., Reviriego, P., Hernández, J. A., Juarez, M., and Sarkar, R. Towards Understanding the Interplay of Generative Artificial Intelligence and the Internet, June 2023.
- Ngufor, C., Van Houten, H., Caffo, B. S., Shah, N. D., and McCoy, R. G. Mixed effect machine learning: A framework for predicting longitudinal change in hemoglobin A1c. *Journal of Biomedical Informatics*, 89:56–67, January 2019. ISSN 1532-0480. doi: 10.1016/j.jbi.2018.09.001.
- Nikolenko, S. I. Synthesis Humans. Create perfectly labeled images & video for ML models. <https://synthesis.ai/synthesis-humans/>, November 2022.
- Noble, S. Algorithms of Oppression: How Search Engines Reinforce Racism. *Science*, 374(6567):542–542, October 2021. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.abm5861.
- Norman, J., Agarwal, S., and Farid, H. An Evaluation of Forensic Facial Recognition, November 2023.
- Park, J., Krishna, R., Khadpe, P., Fei-Fei, L., and Bernstein, M. AI-Based Request Augmentation to Increase Crowdsourcing Participation. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, 7:115–124, October 2019. ISSN 2769-1349. doi: 10.1609/hcomp.v7i1.5282.
- Pasquale, F. Licensure as Data Governance. September 2021.
- Patel, R., Peppin, A., Pavel, V., Brennan, J., Parker, I., and Safak, C. Participatory data stewardship. Technical report, Ada Lovelace Institute, 2021.
- Paullada, A., Raji, I. D., Bender, E. M., Denton, E., and Hanna, A. Data and its (dis)contents: A survey of dataset development and use in machine learning research. *Patterns*, 2(11):100336, November 2021. ISSN 2666-3899. doi: 10.1016/j.patter.2021.100336.

- Peng, K., Mathur, A., and Narayanan, A. Mitigating Dataset Harms Requires Stewardship: Lessons from 1000 Papers, November 2021.
- Perez, D. ChatGPT: Learning Tool — or Threat? How a Texas College Is Eyeing New AI Program, January 2023.
- Qiu, H., Yu, B., Gong, D., Li, Z., Liu, W., and Tao, D. SynFace: Face Recognition With Synthetic Data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10880–10890, 2021.
- Raghavan, M., Barocas, S., Kleinberg, J., and Levy, K. Mitigating bias in algorithmic hiring: Evaluating claims and practices. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT\* '20, pp. 469–481, New York, NY, USA, January 2020. Association for Computing Machinery. ISBN 978-1-4503-6936-7. doi: 10.1145/3351095.3372828.
- Rainie, S. C., Kukutai, T., Walter, M., Figueroa-Rodríguez, O. L., Walker, J., and Axelsson, P. *Indigenous Data Sovereignty*. African Minds and the International Development Research Centre (IDRC), 2019. ISBN 978-1-928331-95-7.
- Raji, I. D. and Fried, G. About Face: A Survey of Facial Recognition Evaluation, February 2021.
- Raji, I. D., Gebru, T., Mitchell, M., Buolamwini, J., Lee, J., and Denton, E. Saving Face: Investigating the Ethical Concerns of Facial Recognition Auditing, January 2020.
- Razmetaeva, Y. The Right to Be Forgotten in the European Perspective. *TalTech Journal of European Studies*, 10(1): 58–76, May 2020. doi: 10.1515/bjes-2020-0004.
- Recht, B., Roelofs, R., Schmidt, L., and Shankar, V. Do ImageNet Classifiers Generalize to ImageNet?, June 2019.
- Robertson, S. and Salehi, N. What If I Don't Like Any Of The Choices? The Limits of Preference Elicitation for Participatory Algorithm Design, July 2020.
- Scheuerman, M. K., Wade, K., Lustig, C., and Brubaker, J. R. How We've Taught Algorithms to See Identity: Constructing Race and Gender in Image Databases for Facial Analysis. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW1):58:1–58:35, May 2020. doi: 10.1145/3392866.
- Scheuerman, M. K., Weathington, K., Mugunthan, T., Denton, E., and Fiesler, C. From Human to Data to Dataset: Mapping the Traceability of Human Subjects in Computer Vision Datasets. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW1):1–33, April 2023. ISSN 2573-0142. doi: 10.1145/3579488.
- Shumailov, I., Shumaylov, Z., Zhao, Y., Gal, Y., Papernot, N., and Anderson, R. The Curse of Recursion: Training on Generated Data Makes Models Forget, May 2023.
- Sloane, M. To make AI fair, here's what we must learn to do. *Nature*, 605:9–9, May 2022. doi: 10.1038/d41586-022-01202-3.
- Sloane, M., Moss, E., Awomolo, O., and Forlano, L. Participation is not a Design Fix for Machine Learning, August 2020.
- Solove, D. J. Privacy Self-Management and the Consent Dilemma. November 2012.
- Solove, D. J. and Hartzog, W. The FTC and the New Common Law of Privacy. *Columbia Law Review*, 114(3): 583–676, 2014.
- Spade, D. *Normal Life: Administrative Violence, Critical Trans Politics, and the Limits of Law*. Duke University Press, July 2015. ISBN 978-0-8223-7479-4.
- Strickler, M. Recent Developments in Privacy Law. *Bus. LAw.*, 76:269, 2020.
- Taboh, J. Facial Recognition Technology Solves Crimes, but at What Cost? <https://www.voanews.com/a/silicon-valley-technology-facial-recognition-technology-solves-crimes-what-cost/6207137.html>, June 2021.
- Thiel, D. Identifying and Eliminating CSAM in Generative ML Training Data and Models. Technical report, Stanford Internet Observatory, December 2023.
- Trinh, T. H., Wu, Y., Le, Q. V., He, H., and Luong, T. Solving olympiad geometry without human demonstrations. *Nature*, 625(7995):476–482, January 2024. ISSN 1476-4687. doi: 10.1038/s41586-023-06747-5.
- Viljoen, S. A Relational Theory of Data Governance. *the yale law journal*, 2021.
- Wang, X. and Huang, H. Approaching Machine Learning Fairness through Adversarial Network, September 2019.
- Wehrli, S., Hertweck, C., Amirian, M., Glüge, S., and Stadelmann, T. Bias, awareness, and ignorance in deep-learning-based face recognition. *AI and Ethics*, 2(3):509–522, August 2022. ISSN 2730-5961. doi: 10.1007/s43681-021-00108-6.
- Whitney, C. D., Naval, T., Quepons, E., Singh, S., Rick, S. R., and Irani, L. HCI Tactics for Politics from Below: Meeting the Challenges of Smart Cities. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–15, Yokohama Japan, May 2021. ACM. ISBN 978-1-4503-8096-6. doi: 10.1145/3411764.3445314.

- Wong, S. C., Gatt, A., Stamatescu, V., and McDonnell, M. D. Understanding data augmentation for classification: When to warp?, November 2016.
- Wood, E., Baltrušaitis, T., Hewitt, C., Dziadzio, S., Cashman, T. J., and Shotton, J. Fake It Till You Make It: Face Analysis in the Wild Using Synthetic Data Alone. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3681–3691, 2021.
- Wu, Z., Yang, Y., Zhao, J., and Wu, Y. The Impact of Algorithmic Price Discrimination on Consumers’ Perceived Betrayal. *Frontiers in Psychology*, 13, 2022. ISSN 1664-1078.
- y Arcas, B. A. Do algorithms reveal sexual orientation or just expose our stereotypes?, January 2018.
- Yang, K., Qinami, K., Fei-Fei, L., Deng, J., and Rusakovsky, O. Towards Fairer Datasets: Filtering and Balancing the Distribution of the People Subtree in the ImageNet Hierarchy. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 547–558, January 2020. doi: 10.1145/3351095.3375709.
- Yew, R.-J. and Xiang, A. Regulating Facial Processing Technologies: Tensions Between Legal and Technical Considerations in the Application of Illinois BIPA. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pp. 1017–1027, Seoul Republic of Korea, June 2022. ACM. ISBN 978-1-4503-9352-2. doi: 10.1145/3531146.3533163.
- Yi, D., Lei, Z., Liao, S., and Li, S. Z. Learning Face Representation from Scratch, November 2014.
- Zhang, L., Wu, Y., and Wu, X. A causal framework for discovering and removing direct and indirect discrimination, November 2016.