# Data Leverage: A Framework for Empowering the Public in its Relationship with Technology Companies

Nicholas Vincent
Northwestern University
nickvincent@u.northwestern.edu

Hanlin Li
Northwestern University
lihanlin@u.northwestern.edu

Nicole Tilly
Northwestern University
nicoletilly2023@u.northwestern.edu

Stevie Chancellor*
University of Minnesota
steviec@umn.edu

Brent Hecht
Northwestern University
bhecht@northwestern.edu

## ABSTRACT

Many powerful computing technologies rely on implicit and explicit data contributions from the public. This dependency suggests a potential source of leverage for the public in its relationship with technology companies: by reducing, stopping, redirecting, or otherwise manipulating data contributions, the public can reduce the effectiveness of many lucrative technologies. In this paper, we synthesize emerging research that seeks to better understand and help people action this *data leverage*. Drawing on prior work in areas including machine learning, human-computer interaction, and fairness and accountability in computing, we present a framework for understanding data leverage that highlights new opportunities to change technology company behavior related to privacy, economic inequality, content moderation and other areas of societal concern. Our framework also points towards ways that policymakers can bolster data leverage as a means of changing the balance of power between the public and tech companies.

## CCS CONCEPTS

• **Human-centered computing → Collaborative and social computing theory, concepts and paradigms**.

## KEYWORDS

data leverage, data strikes, data poisoning, conscious data contribution

**ACM Reference Format:**
Nicholas Vincent, Hanlin Li, Nicole Tilly, Stevie Chancellor, and Brent Hecht. 2021. Data Leverage: A Framework for Empowering the Public in its Relationship with Technology Companies. 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/10.1145/3442188.3445885

*Chancellor completed much of this work while at Northwestern University.


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

ACM ISBN 978-1-4503-8309-7/21/03…$15.00
https://doi.org/10.1145/3442188.3445885

## 1 INTRODUCTION

In August 2020, the most valuable five technology companies had a total market cap of US$7 trillion [98]. This valuation is driven in part by large models that use data generated by the public to recommend content, rank search results, and provide many other services [25, 90, 101, 136]. More generally, lucrative technologies used by many companies rely on data generated by large groups of people to fulfill critical customer needs [8, 25, 101, 124] and drive decision-making [26].

The reliance of powerful technologies (and thus powerful companies) on "data labor" [8, 101, 123] by the general public presents an enormous opportunity for the public to gain more power in its relationship with tech companies. People perform data labor when they engage in the multitude of interactions with technology that generate data for firms (e.g. liking, clicking, rating, posting). By leveraging tech companies' reliance on their data labor, the public could demand changes on pressing issues [60, 73], such as diminished privacy [24, 40], the reinforcement of problematic societal biases by AI systems [5, 7, 44, 73, 95], eroded labor rights [59, 97], environmental harms [108], content moderation challenges [52], and the current imbalance in how profits from data-driven technologies are distributed between tech operators and data contributors [25, 101, 126]. Armed with the knowledge of the importance of data contributions and the tools to action this knowledge, the public could potentially interfere with recommender systems, search engines, image classifiers, and other technologies until tech companies made changes related to these issues.

To capture the power inherent in the public's data labor, this paper introduces the concept of "data leverage" and discusses how the concept can be made operational. Simply put, data leverage refers to influence that members of the public have over tech companies because important computing technologies rely on the public's data contributions. Data leverage catalyzes power achieved by *harming* data-dependent technologies as well as power achieved by *improving* alternative data-dependent technologies and thereby creating increased competition [121]. The concept of data leverage highlights an emergent theme in the FAccT community and related areas, including human-computer interaction (HCI), social computing, society and technology studies (STS), machine learning (ML), and particularly ML research that seeks to advance fairness, justice, and a human-centered perspective (e.g. [5, 20, 29, 54]). This paper shows that this interdisciplinary lens can provide a structure for understanding and actioning an almost entirely untapped source of power that can advance a wide variety of pro-social goals. Our

data leverage framework also highlights opportunities for future research and policy interventions that empower the public in its relationship with technology companies.

The contributions of this work are to (1) define data leverage, (2) provide a framework of potential "data levers", grounded in prior work that has advanced our understanding of these levers, (3) outline an initial assessment of strengths and weaknesses of each data lever in the public's "tool belt", and (4) highlight how data leverage provides important opportunities for research and new policy. Critically, research and policy can amplify data leverage and, conversely, using data leverage as a lens can raise the stakes for related research areas and policy discussions. We pay particular attention to factors that might facilitate the use of data leverage (e.g. policy interventions) or that block groups from exerting data leverage (drawing on the literature on "non-use" of technology) [12–15, 77, 105, 133].

### 1.1 Background and Definitions

Before continuing, we first present formal definitions of data leverage and supporting concepts. Note that while aiming to be comprehensive, these are working definitions. Data leverage is an emerging topic in a rapidly moving field, and we aim to advance and open the discussion around data leverage, not conclude it.

- *Data leverage*: The power derived from computing technologies' dependence on human-generated data. Data leverage is exerted when a group influences an organization by threatening to engage in or directly engaging in data-related actions that harm that organization's technologies or help its competitors' technologies.
- *Data levers*: The specific types of actions that individuals or groups engage in to exert data leverage. For instance, "data strikes" [123] are one of the data levers we discuss below and they operate by cutting off the flow of data to tech companies.

## 2 RELATED WORK

In this section, we situate data leverage in relation to the FAccT domain, and then discuss four additional areas that contribute to the idea of data leverage.

### 2.1 Data Leverage and FAccT Research

Data leverage emerges in part from work in the broader FAccT community that has demonstrated the limitations of purely technical approaches to advancing fairness and justice in computing systems [5, 20, 29, 44, 49]. This large literature emphasizes the critical roles played by the societal context around computing systems, and has demonstrated that sociotechnical approaches are often much more powerful than purely technical approaches. Data leverage can in many ways be understood as a framework that helps us better understand data-driven technologies through a sociotechnical lens and use that lens to take action to achieve pro-social outcomes.

Data leverage is more specifically informed by Kulynych et al.'s work that proposed "Protective Optimization Technologies" (POTs) as a way to address the negative impacts of algorithmic systems and give agency to those impacted [73]. POTs allow people to contest or subvert optimization technologies, perhaps adopting techniques from data poisoning (which we further address below) [73, 119]. Data leverage and POTs are synergistic concepts, and many POTs enable people to exert data leverage.

### 2.2 Data as Labor

Data leverage is heavily informed by work that views data generated by people using computing systems as a type of labor. Building on Posner and Weyl [101], Arrieta Ibarra et al. argue that data should be considered as labor, not "exhaust" emitted in the process of using technology, and as such, should be subject to some kind of remuneration [8]. The relationship between the data-generating public and the companies that benefit from data is very asymmetric. Not only do people have very little knowledge of — let alone agency over — how data they contribute is used, but the economic winnings from powerful data-dependent technologies are reaped entirely by tech companies [101]. To mitigate this inequality, Posner and Weyl called for the formation of "data unions", which allow data laborers to collectively negotiate with technology companies [101].

The discussion around data labor has inspired work that aims to measure the economic value of data [70, 90, 124, 125]. One approach has been to look at the relationship between Wikipedia — the product of data contributions from the public — and real-world economic outcomes such as tourism and investment [62, 134]. Building on the data as labor concept, Vincent and colleagues have investigated how people might withhold or redirect their data labor to force a data-dependent organization to change its practices [121, 123].

Scholars working on data feminism — an intersectional feminism-informed lens for data science — have called for more efforts to make the labor of data science visible, including the labor of data generation [37]. These scholars argue that the invisible labor of data science, much like housework, has been hidden from public view and therefore undervalued [37], and that researchers can begin to shine a light on this labor by studying and highlighting the processes of data creation (e.g. [35]). In this way, data feminism is very aligned with the ideas of data leverage; both aim to measure and make people aware of the value of previously invisible labor and ultimately reshape power imbalances.

### 2.3 Data Leverage and Technology Use/Non-Use

The data leverage concept is also informed by work from HCI and STS on technology "use", "non-use", and the spectrum of behaviors in between.

Work from Selwyn and Wyatt called attention to the need to understand people who do not use new technologies [109, 133]. Most relevant to data leverage, Selwyn documented that people engage in ideological refusal to use certain technology "despite being able to do so in practice". Further calls to study non-use in HCI and STS have been amplified in the years since [11, 105].

Use and non-use exist on a spectrum [15, 133]. People face many social and technical decisions in terms of when they will use, and stop using, a particular technology, and these decisions lead to many different forms of use and non-use [13, 23, 107]. Recently, Saxena et al. reviewed the methods for creating typologies to describe the many forms of use and non-use [106].

Many factors motivate non-use, such as exclusion [133], social capital [77], and socioeconomic factors [12, 15]. Anyone seeking to

use data leverage to empower the public must contend with these factors. Attempts to support data leverage could exclude or disproportionately benefit certain groups following existing patterns in how technology excludes and benefits these groups.

One common theme in the non-use literature is that it is not easy for people to refrain from use when it comes to products that have some benefit in their life, even if the benefit(s) come with a host of long-term drawbacks. People often speak of their technology use as a type of addiction, using terms like 'relapsing" and "withdrawing" [13, 14]. Challenges also emerge related to the public presentation role of social media profiles [76]. Even if people stop using a technology, they may not necessarily delete their data. In studying individuals who left Grindr, a dating app, Brubaker et al. found that "even among those who deleted the app, only a minority tried to close their accounts or remove personal data...[putting them] in a paradoxical position of thinking they have left while their profile — or data — continues on" [23].

The non-use literature also indicates that people engage in protest-related use and non-use behaviors for reasons relating to privacy, data practices, perceived addiction, and other issues [13, 14, 84, 116]. Anyone engaging in such behaviors is a potential participant in data leverage campaigns. Casemajor et al. and Portwood-Stacer argue separately that non-participation in digital media can be an explicitly political action[28, 100]. Li et al. conducted a survey to better understand "protest users", or people who stop or change their use of tech to protest tech companies[84]. The results suggested that there is a large number of people interested in protest use: half of respondents were interested in becoming protest users of a company, and 30% were already engaged in protest use.

An important related lens is that of "refusal". Focusing on bioethics, Benjamin makes the case that broad support of "informed refusal" provides a means of developing a justice-oriented paradigm of science and technology [18]. In practice, people who engage in informed refusal are engaging in a political form of non-use, and thereby data leverage. Building on Benjamin's work, Cifor et al. and Garcia et al. describe how the notion of "critical refusal" informed by feminist scholars can be used improve practices around data [33, 48].

## 2.4 Data Leverage and ML Research

Understanding the full potential of data leverage requires deep engagement with machine learning literature. Two relevant areas of ML research are those that answer questions around (1) the effectiveness of adversarial attacks on data-dependent systems and (2) the relationship between a system's performance and changes to underlying data.

There is a large literature that considers the case of adversaries attempting to attack ML systems (e.g. [10, 19, 30, 45, 56, 75, 78, 81, 92, 99, 112, 113, 115]). In early work on adversarial ML, Barreno et al. developed a taxonomy of attacks on ML systems [10]. They focused in particular on attacks in which an adversary "mis-trains" a system, which is called *data poisoning*. Data poisoning attacks against many types of ML systems have been studied in detail [10, 19, 99, 112, 113, 115]. A type of data poisoning attack that is particularly relevant to the work in this paper is the "shilling" attack, which involves "lying" to a recommender system so that a system recommends certain products favored by the attacker [75]. Accordingly, much work has been done on counteracting shilling (e.g. [30, 56, 78, 92]), which may be of concern to groups who want to use shilling-style data poisoning attacks to exert data leverage as we describe below. Researchers have also explored advanced "data poisoning" techniques that use sophisticated methods to optimally harm ML systems [45, 81], which can be much more effective than unsophisticated attacks (e.g. providing random or average ratings to many items [75]).

Data leverage raises the stakes of the already high-stakes adversarial ML domain. This paper highlights how adversarial techniques, such as data poisoning, are not just relevant to issues of security and privacy, but also to the power dynamics between users and tech companies. While some recent work in adversarial ML has taken a political lens and highlighted real world examples of how adversarial ML can create socially desirable outcomes [7], most of the literature takes a strictly security-oriented lens.

The literature on the relationship between the amount of training data a model has access to and model performance is also highly relevant to data leverage. Many authors have found diminishing returns of additional data across many contexts and algorithms (e.g. [31, 36, 46, 61]), and some have studied techniques to address diminishing returns [21]. These findings are informative as to how effective data leverage can be.

## 2.5 Data Leverage and Data Activism

This paper builds on the literature that explores how the public can change practices of the technology industry. Data activism is a relatively new form of civic participation in response to tech companies' pervasive role in public life [9].

Currently, data activism encompasses practices that affect technology design, development, and deployment [91]. Data leverage can be seen as a subset of data activism with a specific focus on empowering the public to influence the performance of data-dependent technologies. Milan and Van der Velden provided a typology of data activism that further illustrated the specialized activities in this space — proactive and reactive data activism [91]. Proactive data activism refers to activists directly influencing software development or databases through open source projects or collaborating with institutions. A particularly relevant data activism initiative is the open data movement, which aims to democratize information that is currently only accessible to the state or businesses [57]. For example, Baack studied an open data project in Finland and highlighted the intermediary role of data activists between the public and operators of data-dependent technologies [9]. On the other hand, reactive data activism entails activists acting against data-collecting entities through adversarial behaviors such as employing encryption. Data leverage includes both types of data activism.

Equipped with the knowledge and expertise to understand data's role in computing, researchers can provide the public with valuable information to identify and employ effective data leverage practices. Work on data activism has unveiled a rich space to improve data practices [34]. In particular, Lehtiniemi and Ruckenstein called for "linking knowledge production to data activism practice" to gain a comprehensive understanding of data's role in the public sphere [80].

# 3 DATA LEVERAGE FRAMEWORK

In this section, we describe our framework for data leverage in detail . The framework — and this section — is organized around the three data levers we identified. For each lever, we first define the lever and any variants, and do so grounded in past work viewed through our data leverage lens. We then provide practical examples of each data lever and describe the likely factors that will govern the effectiveness of the lever. Table 1 lists the data levers, their definitions, and several examples of each.

## 3.1 Data Strikes

The first of the data levers we will consider are data strikes. Data strikes involve a person withholding or deleting data to reduce the amount of data an organization has available to train and operate data-dependent technologies. Although the term data strike is relatively new, the concept builds on the well-studied practices of stopping or changing technology use as a form of protest, as discussed in Related Work. For instance, groups have participated in prominent boycotts against companies like Facebook and Uber [55, 110]. In another example, people use ad blocking software to deprive companies of data about the success of their ad placements [27].

*3.1.1 Data Strike Variants.* The most basic form of a data strike is a *withholding-based data strike*. In some cases, users can withhold data by reducing or stopping their technology use, or by continuing to use a technology with privacy-protection tools (e.g. tracking blockers [88]). In jurisdictions that allow people to delete their past data (using laws like the General Data Protection Regulation (GDPR) and California Consumer Privacy Act (CCPA) [104, 128]), users can also engage in *deletion-based data strikes*. The effectiveness of such strikes will depend on how well regulations can force companies to regularly retrain or delete their models (so as to remove weights learned using now-deleted data). There is some precedent from the U.S. that model deletion can be enforced: in 2021, the Federal Trade Commission forced a company to delete both customer photos and the facial recognition models trained on the photos [86].

Data strikes can be further categorized based on their coordination requirements. Data strikes (and ther other data levers we will describe below) are likely possible without serious coordination, given the success of hashtag activism [65] and other forms of online collective action that operate without central leadership [87]. For instance, people wanting to start an informally-organized data strike might simply make a call for others to delete as much data as they are willing. However, "targeted" [10] data strikes have the potential for a group of data strikers to achieve disproportionate impact [123]. Following Barreno et al.'s definition of targeted attacks on ML systems, a targeted data strike might encourage participants to delete specific data points or recruit particularly valuable participants. For example, data strikers could try to reduce performance for a specific genre of movie recommendations, while leaving performance for other genres untouched [123]. Leaders might also recruit specific users to join their data strike – power users have disproportionate influence on systems [41, 131, 132] and withholding or deleting their data may be more impactful.

*3.1.2 What Do Data Strikes Look Like in Practice?* To understand what data strikes will look like, we can gain insight from the non-use literature described above. An individual that chooses to use a platform less frequently or avoid a feature of that platform reduces the amount of data they help to generate. In this way, a person's choices about use and non-use affect how much data that person generates. Research in the use and non-use domain has provided empirical examples of what could be conceptualized as data strikes against Facebook and Twitter [12–15, 100, 107].

Privacy and surveillance research also lends itself to uncovering privacy-focused behaviors that can be seen as data strikes. One prominent example is that many people use anti-tracking browser extensions that limit the amount of data online trackers collect [27, 84, 88]. Studies on algorithm transparency also provide evidence suggesting that people engage with data strike-like behaviors because of dissatisfaction with algorithmic system opacity, such as ceasing producing reviews for review platforms [42, 43]. Additionally, research on online communities presented case studies of both Reddit moderators and community members striking by disabling and leaving their communities [89, 94].

*3.1.3 How Can Data Strikes Be Effective?* A data strike can be evaluated based on the importance of the data that "goes missing" in terms of how that data affects relevant data-dependent systems. Said another way, does the missing data noticeably degrade a system's performance, move a classifier's decision boundary (or hyperplane, etc.) in a meaningful way, or otherwise change outputs?

To understand the effectiveness of data strikes, researchers and strike leaders might look to research on data scaling and learning curves, which describes the relationship between ML performance and the amount of data available for training (e.g. [31, 36, 46, 61]). Findings from this literature could be used to predict the effectiveness of a strike, as in prior work which explicitly simulated data strikes [121, 123]. If researchers have shown a model needs a certain number of observations in its training set to be effective (e.g. [31]), data strike organizers could use that research to guide their strike, for instance by setting a goal for participant recruitment.

In summary, data strikes are a data lever available to anyone who can withhold or delete their data. While a new concept, research in HCI, privacy, machine learning, and related fields can help us to understand what data strikes will look like and how effective they might be.

## 3.2 Data Poisoning

A data poisoning attack is an adversarial attack that inserts inaccurate or harmful training data into a data-dependent technology, thereby encouraging the model to perform poorly [10]. While data strikes harm performance by reducing the amount of available data, data poisoning harms performance by providing a technology with data that was created with the intention of thwarting the technology. A relatively accessible way that users can engage in data poisoning is simply by leveraging standard technology features in a deceptive manner. For instance, someone who dislikes pop music might use an online music platform to play a playlist of pop music when they step away from their device with the intention of "tricking" a recommender system into using their data to recommend pop music to similar pop-hating users. Other straightforward examples

**Table 1: The three data levers in our framework, short definitions for each, and several examples of each.**

| Data Lever Name | Short Definition | Examples |
|---|---|---|
| Data Strike | withholding or deleting data | leaving a platform, installing privacy tools |
| Data Poisoning | contributing harmful data | inputting fake data in user profile, clicking randomly, manipulating images |
| Conscious Data Contribution | contributing data to a competitor | switching to a new search engine, transferring photos to a new platform |

include the coordinated effort to create sexually explicit Google search results for former U.S. Senator Rick Santorum's name [51] and coordinated campaigns to use fake reviews to promote certain products [75]. As we will describe below, very sophisticated variants of data poisoning that draw on state-of-the-art machine learning research are also possible.

*3.2.1 Data Poisoning Variants.* Data poisoning is familiar to the ML community through adversarial ML (see e.g [10, 19, 99, 112, 113, 115]) and obfuscation (see e.g. [24, 63]). This means data poisoning organizers can benefit from the knowledge produced through this body of research.

There are many ways an individual alone can engage in data poisoning. The techniques for obfuscation described by Brunton and Nissenbaum are accessible means of data poisoning for individuals. For instance, users might trade accounts (drawing on Brunton and Nissenbaum) or fill in parts of their profile with fake information [32]. As another example, past work has studied attacks that involve following certain Twitter users to throw off Twitter's profiling [93]. These approaches are generally available to an individual acting alone.

The distinction between coordinated data poisoning attacks and uncoordinated attacks is important. Typically, adversarial ML papers frame data poisoning as a contest between a single attacker (which could be an organization) and a defender/victim. In a coordinated data poisoning attack, however, the attacker is an organized collective.

To execute a coordinated data poisoning attack, it will be necessary to find the appropriate technique for a particular technology. Organizers can look to taxonomies in the adversarial ML literature to see what knowledge an attacker requires and what specific systems are vulnerable to attacks [10, 99].

Shilling attacks are a data poisoning variant that focuses on manipulating specific system outcomes rather than general performance degradation [56, 75, 78]. Unlike other poisoning attacks, this type of data leverage manipulates a system to favorably recommend a product that may not actually be high in quality or popularity, i.e. putting "lipstick on a pig". Shilling can be defended against with systems that identify and remove fraudulent or false reviews [85, 96], but these systems themselves may be vulnerable to data poisoning and data strikes.

As with other forms of data leverage, data poisoning applies more generally to any data-dependent technology, not just to ML systems. For instance, Tahmasebian et al. provide a taxonomy of data poisoning attacks against crowdsourcing-based "truth inference" systems [117], e.g. a system that aims to use crowdsourced data to ascertain the true number of cars on a road. Generally, any system that makes or uses estimates about a population can be compromised by sampling poisoned data.

*3.2.2 What Does Data Poisoning Look Like in Practice?* Almost any data-driven technology is vulnerable to deceptive interactions from users, and there are numerous ways to engage in data poisoning in practice. In the wild, there are a wide range of behaviors that constitute data poisoning attacks. Examples include Uber and Lyft drivers providing false information about their availability [79] and internet-browsing user using software to automatically click ads [63].

The most accessible form of data poisoning involves a person using technology in a deceptive manner, e.g. by lying about their personal attributes, watching videos they dislike, or searching for content they are not interested in. They might even use deception-support tools like the location-spoofing software conceptualized by Van Kleek et al. to engage in "computationally-mediated pro-social deception" [120].

By combining findings and tools from HCI and ML, more complex forms of data poisoning may be possible. Users might employ tools like browser extensions (following Li et al. [82] and Howe and Nissenbaum [63]) or web platforms (following Zhang et al. [135]) that help them participate in coordinated data poisoning with sophisticated means of producing poisoned data (e.g. [45, 113]). For instance, one could imagine a data poisoning platform, modeled on existing social computing platforms [72], that provides users with bespoke poisoned data that they can contribute to a data poisoning attack. In such a platform, users could upload images poisoned with pixel-level manipulation to spoof image recognition systems, or take suggestions of content to interact with so as to fool recommender systems.

"Data poisoners" might even take inspiration from recent research on what are known as "adversarial evasion attacks" [113], attacks that help users protect their own images from facial recognition systems (i.e. "evade" the system [99]). Shan et al. show that their tool, Fawkes, can imperceptibly alter images so that state-of-the-art facial recognition cannot recognize the altered images [113]. Such tools might be adapted for data poisoning purposes.

*3.2.3 How can Data Poisoning be Effective?* There are several reasons to believe data poisoning might be a powerful source of data leverage. Recent work on sophisticated data poisoning suggests that very small amounts of poisoned data (e.g. using less than 1% of a training set in work from Geiping et al. [50], using 3% of a training set in work from Steinhardt et al. [115]) can meaningfully change the performance of a classifier. Even unsophisticated data poisoning (e.g. playing music one does not actually enjoy) by a

majority of users could so completely poison a dataset as to make it unusable.

Progress in adversarial ML could actually end up reducing the public's poisoning-based data leverage, in which case non-poisoning data levers would become more important. Fundamentally, data poisoners are engaging in a contest with data scientists. This means any data poisoning technique runs the risk of becoming outdated — if a company's data scientists find or invent a defense, the public might lose leverage [113, 115].

Another interesting outcome of data poisoning is its potential conversion to a data strike. In the case where an organization can detect and delete poisoned data, data poisoning reduces to a data strike. Detectable data poisoning could even be used to replicate a deletion-based data strike. For instance, search engine users could use data poisoning tools such as AdNauseum [63] — which clicks all ads in a user's browser — to effectively make their ad click data useless, forcing the search engine operator to delete it.

In general, to harm a tech company, data poisoning involves deception and requires affecting the experiences of other users of a platform. Consider someone who lies on a dating site, a surprisingly common phenomenon [58, 118]. The user may protect their privacy, but will also poison their own recommendations (e.g. for romantic partners) and make others' dating experiences worse off. The same logic applies to recommendations for friends, videos, and other goods.

A critical challenge for data leverage will be navigating ethical and legal challenges around when data poisoning is acceptable [24, 50, 113, 120]. Whether a particular instance of poisoning is interpreted to be political dissidence or sabotage depends on the society where it is enacted and on case-by-case specifics. For instance, in some cases existing laws around computer abuse or fraud may come into play, such as the United States' Computer Fraud and Abuse Act (CFAA) [2, 64].

## 3.3 Conscious Data Contribution

The above tactics operate by harming, or threatening to a harm, a given data-dependent technology. However, there are cases for which harmful tactics are not a good fit. For instance, perhaps users do not have the regulatory support needed to delete past data [127] or a new technique for detecting poisoned data foils their poisoning attack. Harmful tactics may also be undesirable because an organization's technologies may actively provide benefits to others (e.g. a ML model that is well known to improve accessibility outcomes).

"Conscious data contribution" (CDC) [121] is a promising alternative to harm-based data leverage. In CDC, instead of deleting, withholding, or poisoning data, people give their data to an organization that they support to increase market competition as a source of leverage. People using CDC for data leverage are similar to people engaging in "political consumption" [71], but instead of voting with their wallet, they vote with their data. An exciting aspect of CDC is that while small data strikes struggle to put a dent in large-data technologies because of diminishing returns, CDC by a small group of users takes advantage of diminishing returns and provide a competitor with a large boost in performance. We return to this point later in our assessment of data levers.

*3.3.1 CDC Variants.* Variants of CDC closely mirror variants of data strikes because CDC in a sense is the inverse of data strikes — where data strikes take, CDC gives.

The easiest way to engage in CDC is to simply start using another technology with the intention of producing useful data for the organization that operates the technology. Sometimes, these CDC campaigns may also involve a data strike if a user moves from one platform to another, for example abandoning Google and moving to DuckDuckGo.

In jurisdictions where data portability laws [1] require that companies allow users to download their data, users can engage in CDC by downloading data from a target organization and contributing it to the organization's competitor. Many services already allow users to download or otherwise access some of their data contributions, but the usefulness of currently exportable data to other companies remains to be seen [66].

Similarly to how coordinated data strikes and data poisoning might seek to hurt a particular aspect of a technology, coordinated CDC can enhance specific aspects of a technology's capabilities. In a coordinated CDC campaign, organizers might instruct participants to donate specific types of data, or organizers might seek out specific people to join a campaign, in an effort to focus on contributions towards a specific goal. For instance, in the recommendation context, CDC leaders might seek out comedy movies fans to contribute data to a comedy movie recommender, instead of trying to solicit data about every movie genre. Recommender system researchers have shown that allowing users to filter out their old data could actually improve recommendations [129], so CDC participants could even use filtering to further target their data contributions.

The idea of CDC has complex relationships with various proposals for "data markets" [6, 67], which are designed to give people the ability to sell data that they generate. While data markets allow users to participate in a form of CDC by giving them choices about to whom they will sell data, people may prioritize their personal economic incentives over attempts to gain leverage. A major issue with CDC via data markets is the fact that any data with a social component often has information about more than one person[6, 17], which could make it legally and ethically tricky to handle data via markets.

*3.3.2 What Does CDC Look Like in Practice?* As mentioned above, providing data to online platforms can be a form of Conscious Data Contribution if users aim to increase the performance of these technologies relative to their competitors. As such, there are many existing examples of what CDC might look like in practice.

Cases in which users switch platforms provide one set of examples. In 2015, many Reddit users expressed dissatisfaction with the platform and eventually migrated to alternative platforms such as Voat and Snapzu [94]. In doing so, these users performed an act of CDC, explicitly supporting Reddit's competitors. Past work suggests that migrations are an especially likely form of CDC, because an individual user's choice to move platforms as part of a CDC campaign may lead to people in the user's social network also migrating [47, 74]. Where social networks create friction against data strikes, they can help to drive CDC.

Many research initiatives involve collecting volunteered data, which in certain cases could provide opportunities for CDC. In Silva et al.'s study, people contributed data about their Facebook political ads to researchers for monitoring and auditing purposes [114]. While research studies on their own are not necessarily CDC (though they could be, if the research helps support competitive data-driven technologies), they can often provide a good example of how CDC might be implemented.

Other types of data sharing and generation can also be CDC. For instance, the "data donation" concept explored in the context of data ethics [102] could be used for CDC. In some cases participation in human computation [103], crowdsourcing systems, and other social computing platforms [72] could qualify as CDC. For example, under our definition, people who choose to contribute data to protein-folding games could be engaging in a form of CDC [39], with the potential to exert leverage against other organizations that benefit from protein folding models.

*3.3.3 How Can CDC Be Most Effective?* CDC has a lower barrier to entry than data strikes and data poisoning because it is possible to engage in CDC without completely stopping use of an existing technology. Despite this advantage, a critical question for any CDC effort will be how much leverage "helping a competitor" exerts on the target. For instance, a group of CDC users might be able to successfully improve the ML technologies of a small startup that is competing with a major platform. However, even with improved data-driven technologies, other factors like access to capital and switching costs for users might prevent the startup from competing effectively with the original target of leverage, thus reducing the chance that the original target changes their practices. In some cases, standing up a viable competitor that has better practices could be the end goal of a CDC campaign, even if does not directly harm another company. By supporting a new viable contender, CDC participants can effectively change the overall relationship between the public and technology companies.

Like data strikes, a key determinant of the effectiveness of CDC will be the level of participation. The more people that participate in CDC, the more powerful it will become, and the degree of effectiveness can be estimated using ML findings and methods as we discuss below. A critical distinction between data strikes and CDC is that while small data strikes may struggle to escape the flat portion of ML learning curves, CDC by a small group can actually provide a huge boost in ML performance to a small organization. We expand on this comparison in the following Assessment section.

# 4 ASSESSING DATA LEVERS

In this section, we use three axes to evaluate strengths and weaknesses of each data lever: the *barrier-to-entry* to use a data lever, how *ethical and legal considerations* might complicate the use of a data lever, and finally the *potential impact* of each data lever. Table 2 contains a brief summary of our assessments.

## 4.1 Barriers to Entry

In general, CDC has the lowest barrier to entry of the data levers we identified. This is because CDC does not require stopping or changing the use of existing technologies, which prior work discussed above indicates can be challenging (e.g. [13, 14, 107]). A person can continue using existing technologies operated by an organization against which they want to exert leverage while engaging in CDC [69, 121]. The main barriers to transfer-based CDC are regulatory and technical. Do laws help people transfer their data [1] and do tools exist to make data transfer realistic?

The barriers to entry for data strikes are more substantial then those for CDC and less substantial than those for data poisoning. While participating in a data strike disrupts a user's access to online platforms, strikes do not necessarily force a user to stop using a platform like a traditional boycott would. For instance, a user who relies on Facebook to communicate with family members could stop engaging with sponsored content on Facebook but continue messaging their family members. An Amazon user might continue buying products but stop leaving ratings and reviews. An important downside of data strikes is that they hurt the performance of technologies for participating users. By cutting off data contributions, an individual often reduces their own ability to benefit from a system. As discussed by Vincent et al. [123], the effect of a data strike will almost always be most pronounced on the strike participants.

The barriers to entry for each data lever are also contingent on the bandwidth available to potential participants and any potential data caps or data charges they have. Data strikes are likely the least limited by bandwidth (although striking against an Internet provider, e.g. Facebook Free Basics, could be challenging [111]). In places where the Internet is easy to access and has relatively high data caps, poisoning data by letting music stream for hours or actively manipulating multimedia may be accessible. In contrast, in places where Internet access is limited [38], poisoning data may be difficult if not impossible. Similar dynamics likely will apply to CDC: data caps could stifle efforts to engage in CDC.

Many of the barriers to entry discussed above are not equally distributed across different populations, and this means that different populations likely have differing access to data leverage. For instance, with regards to data poisoning, the time available to expend the necessary effort and/or the skills necessary to do so will limit the ability of many populations to engage in data poisoning. Those most positioned to perform data poisoning attacks are ML researchers, technologists, and others with strong technical skills, an already relatively privileged group. Nonetheless, members of this group could use their powerful position for the benefit of people without these advantages (there is precedent of tech worker organizing along these lines [3]).

Turning to coordination, data leverage campaigns will differ in their coordination needs, with greater coordination requirements raising the barrier to entry for all three data levers. Large-scale data leverage is possible without formal organization: boycotts using Twitter hashtags provide real-world examples [65]. However, certain data levers require especially well-coordinated effort to see impact, e.g. sophisticated data poisoning [45].

## 4.2 Legal and Ethical Considerations

Data leverage organizers may face legal and ethical challenges. Withholding-based data strikes face the fewest of these challenges. These data strikes require almost no regulatory support as users can simply cease using platforms (keeping in mind the differential barrier to entry concerns discussed above). Deletion-based data

**Table 2: Summary of key points from our assessment of data levers.**

| Data Lever | Barriers to Entry | Legal and Ethical Considerations | Potential Impact |
|---|---|---|---|
| Data Strike | *moderate*:<br>-non-use is challenging<br>-hurts participating users<br>-need for privacy tools | *lower*:<br>-need privacy laws to delete data<br>-harming tech may be undesirable | *moderate*:<br>-small group has small effect<br>-large group can have huge impact |
| Data Poisoning | *higher*:<br>-time/effort/bandwidth costs<br>-may require ML knowledge<br>-may require extra coordination | *higher*:<br>-potentially illegal<br>-harming tech may be undesirable<br>-inherently deceptive | *moderate*:<br>-small group can have huge effects<br>-if caught, "reduces" to a strike<br>-constant arms race |
| Conscious Data Contribution | *lower*:<br>-can continue using existing tech | *moderate*:<br>-potential to improve harmful technologies<br>-privacy concerns of sharing data | *moderate*:<br>-small group can have large effects<br>-large group faces diminishing returns |

strikes require a right to deletion and a guarantee that companies are not *data laundering* by retaining model weights trained on old data [86].

The legality of data poisoning is likely to remain an open question, and interdisciplinary work between computer scientists and legal scholars will be critical to understand the legal viability of data poisoning as a type of data leverage (and to do so in different jurisdictions). Arguments about the ethics of obfuscation (which itself can be a form of data poisoning) raised by Brunton and Nissenbaum apply directly to the use of all types of data poisoning [24]. Participants must contend with the potential effects of dishonesty, wastefulness, and other downstream effects of data poisoning. For instance, there are many harms that could stem from poisoning systems that improve accessibility, block hate speech, or support medical decision-making.

Interesting legal and ethical questions also emerge around CDC. Notably, if a certain data-driven technology is fundamentally harmful and no version of it can meaningful reduce harms (as can be argued for e.g. certain uses of facial recognition [5, 49]), CDC will effectively be neutralized.

Another challenge specific to CDC is that there is the potential that data contributions by one person might violate the privacy of others, as data is rarely truly "individual" [6, 17]. For instance, genetic data about one individual may reveal attributes about their family, while financial data may reveal attributes about their friends. On the legal front, CDC often requires either regulatory support in the form of data portability laws or data export features from tech companies.

## 4.3 Potential Impact

Data strikes and data poisoning harm data-dependent technologies, while CDC improves the performance of a data-dependent technology that can then compete with the technology that is the target of data leverage. We can measure potential impact in terms of performance improvement/degradation, as well as downstream effects (e.g. performance degradation leads to users leaving a platform). Ultimately, we are interested in how likely a data lever is to successfully change an organization's behavior with regards to the goals of the data leverage effort, e.g. making changes related to economic inequality, privacy, environmental impact, technologies that reinforce bias, etc.

A relevant finding from prior work [126] describes how data strikes interact with diminishing returns of data. ML performance exhibits diminishing returns; in general, for a particular task, a system can only get so accurate even with massive increases in available data. As such, when an organization accumulates a sufficient amount of data and begins to receive diminishing returns from new data, that organization is not very vulnerable to small data strikes. Such strikes will — broadly speaking — only unwind these diminishing marginal returns. To a company with billions of users, a (relatively) small data strike simply may not matter.

The potential impact of data poisoning is also enormous: a large-scale data poisoning attack could render a dataset completely unusable. This approach is also appealing for bargaining: a group could poison some data contributions, and make some demand in return for the "antidote". However, the enormous corporate interest in detecting data poisoning means that the would-be poisoners face a constant arms race with operators of targeted technologies. In the worst case scenario, they will be caught, their poisoned data deleted, and the end effect will be equivalent to a data strike.

CDC campaigns, which improve technology performance, operate in the opposite direction of data strikes. Small-scale CDC could be high impact: about 20% of the users of a system could help a competitor get around 80% of the best-case performance [121]. On the other hand, once returns begin to diminish, the marginal effect of additional people engaging in CDC begins to fall.

Given the current evidence, we believe that the data levers we described have a place in the tool belt of those seeking to change the relationship between tech companies and the public. A critical challenge for data leverage researchers will be identifying the correct tool for a specific job. Based on the technologies a target organization uses, a realistic estimate of how many people might participate in data leverage, and knowledge about the resources available to participants, which data lever is most effective?

# 5 DISCUSSION

In this section, we discuss questions associated with data leverage that lie beyond the bounds of our current framework. We first discuss the key question of who might expect to benefit from data

leverage, and highlight how data leverage might backfire. Next, we summarize key opportunities for researchers, particularly those working in or around FAccT topics. Finally, we summarize opportunities for policy that can amplify and augment data leverage.

## 5.1 Who Benefits from Data Leverage?

Researchers, practitioners, activists, policymakers and others interested in studying, supporting, or amplifying data leverage to reduce power imbalances must contend with unequal access to data leverage. As discussed above, there is strong reason to expect that inequalities in access to data leverage mirror known patterns in access to technology and other sources of power more generally [5]. However, our framework suggests that data poisoning and CDC in particular might allow small groups to have disproportionate impacts. A group of users with needs not currently met by existing technologies might engage in CDC to support a competitor to existing tech companies, or use sophisticated data poisoning techniques that require coordination and knowledge, but not mass participation. Researchers can play an active and critical role by developing tools and promoting policy that widely distributes the ability to participate in data leverage efforts and receive benefits from data leverage. Future work may also need to contend with the possibility of organizations counteracting data leverage, e.g. removing access to publicly available data to maintain a dominant market position.

## 5.2 Data Leverage and Data in the Commons

Many lucrative data-dependent technologies rely on "commons" data (e.g. Wikipedia and OpenStreetMap) in addition to the largely proprietary types of data we have discussed so far (e.g. interaction data, rating data). The same is largely true for a variety of data sources that are privately-owned but are a sort of de facto commons for many purposes (e.g. Reddit data, public Twitter posts). Examples of commons-dependent technologies include large language models (e.g. [22]), search engines (e.g. [90, 122, 125]), and a variety of geographic technologies (e.g. [68]). Commons datasets have also been instrumental to the advancement of ML research (e.g.[4, 16]).

How can we view the widespread dependence on commons datasets through the lens of data leverage? Adopting a narrow perspective, all three data levers can certainly be employed using data in the commons. In fact, doing so might be a very effective way of exerting data leverage against a large number of data-dependent technologies at once. For instance, through poisoning (i.e. vandalizing) Wikipedia, one can negatively affect a wide variety of Wikipedia-dependent technologies including Google Search, Bing, and Siri [90, 122, 125]. Indeed, this has already been done with humorous intent a number of times (e.g. [130]). One could similarly imagine organizing a "data strike" of sorts in Wikipedia or OpenStreetMap that sought to ensure that a certain type of information does not appear in these datasets.

That said, from a broader perspective, it is very likely that data poisoning and data strikes using commons data will do substantially more harm than good. For instance, a concerted effort to vandalize (i.e. poison) Wikipedia will cause substantial damage: it would harm Wikipedia readers across the world and would affect technologies operated by non-targeted organizations in addition to those operated by targeted ones. A similar case could be made for most data strikes.

CDC in the context of commons data presents a more complex set of considerations. Indeed, contributing to a commons dataset like Wikipedia can in some ways be understood as a type of CDC as it helps smaller organizations as well as larger ones. However, an important consideration here is that the ability to make use of commons datasets in data-driven technologies is gated by capital. A salient example is GPT-3, OpenAI's high-profile language model that uses training data from sources like Wikipedia and Reddit [22]. The unprecedented computing power needed to train GPT-3 highlights how the data labor that improves Wikipedia and Reddit can disproportionately benefit organizations with enormous resources. An unfortunate reinforcing dynamic regarding data leverage and commons data thus emerges: while a huge number of organizations and individuals stand to be harmed by any sort of poisoning attack or strike on commons data, large and wealthy firms often stand to benefit disproportionately from improvements to these data. Future work that focuses on efficient training, smaller models, and related goals can help to mitigate this particular concern. Similarly, efforts to open-source models themselves (e.g. share model weights) could also help.

## 5.3 Can Data Leverage Research Backfire?

We have presented data leverage as a means to empower the public to address concerns around computing systems that exacerbate power imbalances and create negative societal outcomes. However, research, tools, and policy intended to help data leverage achieve these goals could do the opposite by empowering groups to perpetuate inequalities and, therefore, achieve socially harmful outcomes. For instance, hate groups take advantage of "data voids" in search engines to engage in what can be understood as data poisoning attacks by inserting hateful content and influencing model development [53]. Why wouldn't these groups also try to use other types of data leverage for similar ends?

There are no clear-cut ways to eliminate these risks, but there are steps that data leverage researchers can take to avoid a "backfire" outcome. When designing tools to support data leverage, designers might consider heuristic preventative design from Li et al. [82] and try to make harmful uses of a technology more challenging. For instance, a data poisoning tool might only help users poison certain types of images known to be important to a particular company or technology. Designers should also consider the principles of data feminism [33, 48], including those that emphasize challenging existing hierarchies, embracing pluralism and context, and making labor visible.

## 5.4 Key Research Opportunities for Data Leverage

The concept of data leverage presents exciting research opportunities for many fields. Researchers in FAccT, ML, HCI, STS and related areas in particular have unique opportunities to amplify data leverage.

Data leverage presents a new way of exerting pressure on corporations to make important changes. Most relevant to the FAccT

community, this might involve exerting leverage so that a tech company stops the use of a harmful algorithm [73], or pushing for new economic relationships between data contributors and AI operators in which the benefits of AI are shared more broadly [101, 126]. Data leverage thus presents a novel avenue for researchers to actively pursue pro-social research roles and goals [5].

There is enormous potential to support data leverage with ML research methods. Using simulations and small-scale experiments, future work could build a catalog of results that activists could draw on to make predictions about the effectiveness of a particular data lever in a particular context, such as "if we get $x$ participants to engage in a data strike against technology $y$, we can expect to bring down the accuracy of technology by $z$%, which will likely be enough to encourage company $c$ to make the changes we are demanding". As data leverage becomes more mainstream, there may also be opportunities to study real-world examples and answer key questions such as: What are the downstream effects on revenue, user retention, and actual changes in company behavior?

Future design work could build upon the collective action literature and develop tools to coordinate efforts to use data leverage. For example, because collective action's progress is often opaque to individual participants and this can negatively impact engagement, future work may adopt tactics from "boycott-assisting technologies" [82] and display the impact of the public's data strike or poisoning (e.g. this technology has lost 3% of data). Such tools could also support automating data strikes or data poisoning, similar to AdNauseam Howe and Nissenbaum, to lower the barrier to entry for the public.

In addition to data strikes and poisoning, the computing community can also support CDC by addressing data compatibility and portability issues across platforms and technologies. Data generated by users are often highly platform- and/or technology-dependent. For example, ratings for the same restaurant or hotel may vary significantly across review platforms [42, 83]. Directly transferring data from one technology to another as an act of CDC may run into compatibility issues and even negatively affect the recipient's performance. There is a need for researchers and practitioners to develop software that automatically translates data generated using one technology into data can truly benefit another technology to maximize the success of CDC-based approaches.

Researchers should also seek to better understand the full set of societal impacts that would result from the widespread use of data leverage. As we have discussed above, we hypothesize that the direct effects of actioning data leverage will often involve broadly positive societal impacts, e.g. improved privacy, better distribution of the economic benefits from AI systems, more democratic governance of AI systems. However, the second- and greater-order effects of these changes are more difficult to assess, and even some direct effects may be negative in some cases as highlighted previously. More generally, data leverage defines a pathway to altering power structures in the current computing paradigm. Alterations of power structures in such a complex sociotechnical environment will almost certainly lead to complex outcomes, and more research will be needed to understand these potential outcomes.

### 5.5 Key Policy Opportunities for Data Leverage

Data leverage stands to benefit heavily from regulatory support. As such, data leverage research should be deeply engaged with policy by highlighting regulatory approaches that are likely to amplify the power of data leverage and address its potential negative impacts. Our taxonomy only scratches the surface of how policy may support data leverage; we are excited for this important direction of future work.

Following directly from our assessment of data levers above, we suggest a variety of ways policy can support data leverage:

- Data portability laws will directly enhance CDC, enabling users to contribute data they helped generate in the past.
- Right-to-delete laws will enhance data strikes, assuming these laws also account for the possibility that companies might "launder" deleted data in model weights.
- Data transparency laws that make data collection more apparent may help foster support for data leverage movements.

We note that these policy suggestions are generally aligned with policy aimed at addressing privacy concerns. This suggests a potential "win-win" situation, in which policy simultaneously supports consumer privacy and enhances data leverage.

Expanding on the above points about data portability and right-to-delete laws, policy also offers the potential for making it easy for individuals to use multiple data levers in conjunction with one another. As mentioned above, there are natural connections between data strikes and CDC: by moving from one platform to a new platform, a user can take advantage of both data levers. However, through regulatory support, it may be possible to engage in much more elaborate combinations of data strikes and CDC, for instance deleting only certain pieces of data and transferring over other pieces of data.

## 6 CONCLUSION

In this paper, we presented a framework for using "data leverage" to give the public more influence over technology company behavior. Drawing on a variety of research areas, we described and assessed the "data levers" available to the public. We highlighted key areas where researchers and policymakers can amplify data leverage and work to ensure data leverage distributes power more broadly than is the case in the status quo.

## 7 ACKNOWLEDGMENTS

This work was funded in part by NSF grants 1815507 and 1707296. We are grateful for feedback from colleagues at the CollabLab at Northwestern, GroupLens at the University of Minnesota, and the Community Data Science Collective.

## REFERENCES

[1] 2018. Art. 20 GDPR – Right to data portability | General Data Protection Regulation (GDPR). https://gdpr-info.eu/art-20-gdpr
[2] 2020. 18 U.S. Code § 1030 - Fraud and related activity in connection with computers. https://www.law.cornell.edu/uscode/text/18/1030 [Online; accessed 7. Oct. 2020].
[3] 2020. Tech Workers Coalition. https://techworkerscoalition.org [Online; accessed 6. Oct. 2020].
[4] 2020. Wikipedia:Academic studies of Wikipedia - Wikipedia. https://en.wikipedia.org/w/index.php?title=Wikipedia:Academic_studies_of_Wikipedia&oldid=971074694 [Online; accessed 29. Sep. 2020].

[5] 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.

[6] Daron Acemoglu, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar. 2019. *Too much data: Prices and inefficiencies in data markets*. Technical Report. National Bureau of Economic Research.

[7] Kendra Albert, Jon Penney, Bruce Schneier, and Ram Shankar Siva Kumar. 2020. Politics of Adversarial Machine Learning. In *Towards Trustworthy ML: Rethinking Security and Privacy for ML Workshop, Eighth International Conference on Learning Representations (ICLR)*.

[8] Imanol Arrieta Ibarra, Leonard Goff, Diego Jiménez Hernández, Jaron Lanier, and E Weyl. 2018. Should We Treat Data as Labor? Moving Beyond 'Free'. *American Economic Association Papers & Proceedings* 1, 1 (2018).

[9] Stefan Baack. 2015. Datafication and empowerment: How the open data movement re-articulates notions of democracy, participation, and journalism. *Big Data & Society* 2, 2 (2015), 2053951715594634. https://doi.org/10.1177/2053951715594634 arXiv:https://doi.org/10.1177/2053951715594634

[10] Marco Barreno, Blaine Nelson, Russell Sears, Anthony D Joseph, and J Doug Tygar. 2006. Can machine learning be secure?. In *Proceedings of the 2006 ACM Symposium on Information, computer and communications security*. 16–25.

[11] Ames Morgan G. Brubaker Jed R. Burrell Jenna Dourish Paul Baumer, Eric PS. 2014. Refusing, Limiting, Departing: Why We Should Study Technology Non-Use. In *CHI EA '14: CHI '14 Extended Abstracts on Human Factors in Computing Systems*. 65–68.

[12] Eric PS Baumer. 2018. Socioeconomic Inequalities in the Non use of Facebook. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–14.

[13] Eric PS Baumer, Phil Adams, Vera D Khovanskaya, Tony C Liao, Madeline E Smith, Victoria Schwanda Sosik, and Kaiton Williams. 2013. Limiting, leaving, and (re) lapsing: an exploration of facebook non-use practices and experiences. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 3257–3266.

[14] Eric PS Baumer, Shion Guha, Emily Quan, David Mimno, and Geri K Gay. 2015. Missing photos, suffering withdrawal, or finding freedom? How experiences of social media non-use influence the likelihood of reversion. *Social Media+ Society* 1, 2 (2015), 2056305115614851.

[15] Eric PS Baumer, Shion Guha, Patrick Skeba, and Geraldine Gay. 2019. All Users are (Not) Created Equal: Predictors Vary for Different Forms of Facebook Non/use. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–28.

[16] Jason Baumgartner, Savvas Zannettou, Brian Keegan, Megan Squire, and Jeremy Blackburn. 2020. The pushshift reddit dataset. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14. 830–839.

[17] Omri Ben-Shahar. 2019. Data Pollution. *Journal of Legal Analysis* 11 (2019), 104–159. Publisher: Narnia.

[18] Ruha Benjamin. 2016. Informed refusal: Toward a justice-based bioethics. *Science, Technology, & Human Values* 41, 6 (2016), 967–990. Publisher: SAGE Publications Sage CA: Los Angeles, CA.

[19] Battista Biggio, Blaine Nelson, and Pavel Laskov. 2012. Poisoning attacks against support vector machines. *arXiv preprint arXiv:1206.6389* (2012).

[20] Reuben Binns. 2018. Fairness in machine learning: Lessons from political philosophy. In *Conference on Fairness, Accountability and Transparency*. 149–159.

[21] Michael Bloodgood and Chris Callison-Burch. 2010. Bucking the trend: large-scale cost-focused active learning for statistical machine translation. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 854–864.

[22] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. arXiv:cs.CL/2005.14165

[23] Jed R Brubaker, Mike Ananny, and Kate Crawford. 2016. Departing glances: A sociotechnical account of 'leaving'Grindr. *New Media & Society* 18, 3 (2016), 373–390.

[24] Finn Brunton and Helen Fay Nissenbaum. 2015. *Obfuscation: a user's guide for privacy and protest*. MIT Press, Cambridge, Massachusetts.

[25] Erik Brynjolfsson and Andrew McAfee. 2014. *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.

[26] Erik Brynjolfsson and Kristina McElheran. 2016. The rapid adoption of data-driven decision-making. *American Economic Review* 106, 5 (2016), 133–39.

[27] Ceren Budak, Sharad Goel, Justin Rao, and Georgios Zervas. 2016. Understanding Emerging Threats to Online Advertising. In *Proceedings of the 2016 ACM Conference on Economics and Computation (EC '16)*. ACM, New York, NY, USA, 561–578. https://doi.org/10.1145/2940716.2940787 event-place: Maastricht, The Netherlands.

[28] Nathalie Casemajor, Stıfmmode\acutee\elseé\fiphane Couture, Mauricio Delfin, Matthew Goerzen, and Alessandro Delfanti. 2015. Non-participation in digital media: toward a framework of mediated political action. *Media, Culture & Society* 37, 6 (May 2015), 850–866. https://doi.org/10.1177/0163443715584098 Publisher: SAGE Publications Ltd.

[29] Stevie Chancellor, Eric PS Baumer, and Munmun De Choudhury. 2019. Who is the" Human" in Human-Centered Machine Learning: The Case of Predicting Mental Health from Social Media. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–32.

[30] Paul-Alexandru Chirita, Wolfgang Nejdl, and Cristian Zamfir. 2005. Preventing shilling attacks in online recommender systems. In *Proceedings of the 7th annual ACM international workshop on Web information and data management*. 67–74.

[31] Junghwan Cho, Kyewook Lee, Ellie Shin, Garry Choy, and Synho Do. 2015. How much data is needed to train a medical image deep learning system to achieve necessary high accuracy? *arXiv preprint arXiv:1511.06348* (2015).

[32] Max Cho. 2011. Unsell Yourself—A Protest Model Against Facebook. *Yale Law & Technology* (2011).

[33] Marika Cifor, Patricia Garcia, TL Cowan, Jasmine Rault, Tonia Sutherland, Anita Say Chan, Jennifer Rode, Anna Lauren Hoffmann, Niloufar Salehi, and Lisa Nakamura. 2019. Feminist data manifest-no.

[34] Nick Couldry and Alison Powell. 2014. Big data from the bottom up. *Big Data & Society* 1, 2 (2014), 2053951714539277.

[35] Kate Crawford and Vladan Joler. 2018. Anatomy of an AI System-The Amazon Echo as an anatomical map of human labor, data and planetary resources. *AI Now Institute and Share Lab* 7 (2018).

[36] Thomas G Dietterich and Eun Bae Kong. 1995. *Machine learning bias, statistical bias, and statistical variance of decision tree algorithms*. Technical Report. Technical report, Department of Computer Science, Oregon State University.

[37] Catherine D'Ignazio and Lauren F Klein. 2020. *Data feminism*. MIT Press.

[38] Michaelanne Dye, David Nemer, Laura R Pina, Nithya Sambasivan, Amy S Bruckman, and Neha Kumar. 2017. Locating the Internet in the Parks of Havana. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 3867–3878.

[39] Christopher B Eiben, Justin B Siegel, Jacob B Bale, Seth Cooper, Firas Khatib, Betty W Shen, Barry L Stoddard, Zoran Popovic, and David Baker. 2012. Increased Diels-Alderase activity through backbone remodeling guided by Foldit players. *Nature biotechnology* 30, 2 (2012), 190–192.

[40] Michael D. Ekstrand, Rezvan Joshaghani, and Hoda Mehrpouyan. 2018. Privacy for All: Ensuring Fair and Equitable Privacy Protections *(Proceedings of Machine Learning Research)*, Sorelle A. Friedler and Christo Wilson (Eds.), Vol. 81. PMLR, New York, NY, USA, 35–47. http://proceedings.mlr.press/v81/ekstrand18a.html

[41] Farzad Eskandanian, Nasim Sonboli, and Bamshad Mobasher. 2019. Power of the Few: Analyzing the Impact of Influential Users in Collaborative Recommender Systems. In *Proceedings of the 27th ACM Conference on User Modeling, Adaptation and Personalization*. 225–233.

[42] Motahhare Eslami, Kristen Vaccaro, Karrie Karahalios, and Kevin Hamilton. 2017. " Be Careful; Things Can Be Worse than They Appear": Understanding Biased Algorithms and Users' Behavior Around Them in Rating Platforms.. In *ICWSM*. 62–71.

[43] Motahhare Eslami, Kristen Vaccaro, Min Kyung Lee, Amit Elazari Bar On, Eric Gilbert, and Karrie Karahalios. 2019. User attitudes towards algorithmic opacity and transparency in online reviewing platforms. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14.

[44] Virginia Eubanks. 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.

[45] Minghong Fang, Neil Zhenqiang Gong, and Jia Liu. 2020. *Influence Function based Data Poisoning Attacks to Top-N Recommender Systems*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3366423.3380072

[46] Rosa L Figueroa, Qing Zeng-Treitler, Sasikiran Kandula, and Long H Ngo. 2012. Predicting sample size required for classification performance. *BMC medical informatics and decision making* 12, 1 (2012), 8. https://link.springer.com/article/10.1186/1472-6947-12-8 Publisher: Springer.

[47] David Garcia, Pavlin Mavrodiev, and Frank Schweitzer. 2013. Social resilience in online communities: The autopsy of friendster. In *Proceedings of the first ACM conference on Online social networks*. 39–50.

[48] Patricia Garcia, Tonia Sutherland, Marika Cifor, Anita Say Chan, Lauren Klein, Catherine D'Ignazio, and Niloufar Salehi. 2020. No: Critical Refusal as Feminist Data Practice. In *Conference Companion Publication of the 2020 on Computer Supported Cooperative Work and Social Computing*. 199–202.

[49] Timnit Gebru. 2019. Oxford Handbook on AI Ethics Book Chapter on Race and Gender. *arXiv preprint arXiv:1908.06165* (2019).

[50] Jonas Geiping, Liam Fowl, W. Ronny Huang, Wojciech Czaja, Gavin Taylor, Michael Moeller, and Tom Goldstein. 2020. Witches' Brew: Industrial Scale Data Poisoning via Gradient Matching. arXiv:cs.CV/2009.02276

[51] Tarleton Gillespie. 2017. Algorithmically recognizable: Santorum's Google problem, and Google's Santorum problem. *Information, communication & society* 20, 1 (2017), 63–80.

[52] Tarleton Gillespie. 2018. *Custodians of the Internet: Platforms, content moderation, and the hidden decisions that shape social media.* Yale University Press.

[53] Michael Golebiewski and Danah Boyd. 2019. Data voids: Where missing data can easily be exploited. *Data & Society* (2019).

[54] Ben Green. 2018. 'Fair' Risk Assessments: A Precarious Approach for Criminal Justice Reform. In *5th Workshop on fairness, accountability, and transparency in machine learning*.

[55] Rebecca Greenfield, Sarah Frier, and Ben Brody. 2018. NAACP Seeks Week-Long Facebook Boycott Over Racial Targeting. *Bloomberg.com* (Dec. 2018). https://www.bloomberg.com/news/articles/2018-12-17/naacp-calls-for-week-long-facebook-boycott-over-racial-targeting

[56] Ihsan Gunes, Cihan Kaleli, Alper Bilge, and Huseyin Polat. 2014. Shilling attacks against recommender systems: a comprehensive survey. *Artificial Intelligence Review* 42, 4 (2014), 767–799. Publisher: Springer.

[57] Michael Gurstein. 2011. Open data: Empowering the empowered or effective data use for everyone? *First Monday* 16 (02 2011). https://doi.org/10.5210/fm.v16i2.3316

[58] Jeffrey T Hancock, Catalina Toma, and Nicole Ellison. 2007. The truth about lying in online dating profiles. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 449–452.

[59] Kotaro Hara, Abigail Adams, Kristy Milland, Saiph Savage, Chris Callison-Burch, and Jeffrey P Bigham. 2018. A Data-Driven Analysis of Workers' Earnings on Amazon Mechanical Turk. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 449.

[60] B Hecht, L Wilcox, JP Bigham, J Schöning, E Hoque, J Ernst, Y Bisk, L De Russis, L Yarosh, B Anjum, and others. 2018. It's time to do something: Mitigating the negative impacts of computing through a change to the peer review process.

[61] Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md Patwary, Mostofa Ali, Yang Yang, and Yanqi Zhou. 2017. Deep learning scaling is predictable, empirically. *arXiv preprint arXiv:1712.00409* (2017).

[62] Marit Hinnosaar, Toomas Hinnosaar, Michael E Kummer, and Olga Slivko. 2019. Wikipedia matters. *Available at SSRN 3046400* (2019).

[63] Daniel C Howe and Helen Nissenbaum. 2017. Engineering Privacy and Protest: A Case Study of AdNauseam.. In *IWPE@ SP*. 57–64.

[64] Kate Mathews Hunt. 2015. Gaming the system: Fake online reviews v. consumer law. *Computer law & security review* 31, 1 (2015), 3–25.

[65] Sarah J Jackson, Moya Bailey, and Brooke Foucault Welles. 2020. *# HashtagActivism: Networks of Race and Gender Justice.* MIT Press.

[66] Ross James. 2020. How to use Google Takeout to download your Google data - Business Insider. *Business Insider* (Jan 2020). https://www.businessinsider.com/what-is-google-takeout

[67] Ruoxi Jia, David Dao, Boxin Wang, Frances Ann Hubis, Nick Hynes, Nezihe Merve Gürel, Bo Li, Ce Zhang, Dawn Song, and Costas J Spanos. 2019. Towards Efficient Data Valuation Based on the Shapley Value. In *The 22nd International Conference on Artificial Intelligence and Statistics*. 1167–1176.

[68] Isaac L. Johnson, Yilun Lin, Toby Jia-Jun Li, Andrew Hall, Aaron Halfaker, Johannes Schöning, and Brent Hecht. 2016. *Not at Home on the Range: Peer Production and the Urban/Rural Divide.* Association for Computing Machinery, New York, NY, USA, 13–25. https://doi.org/10.1145/2858036.2858123

[69] Charles I Jones and Christopher Tonetti. 2019. *Nonrivalry and the Economics of Data.* Technical Report. National Bureau of Economic Research.

[70] Pang Wei W Koh, Kai-Siang Ang, Hubert Teo, and Percy S Liang. 2019. On the accuracy of influence functions for measuring group effects. In *Advances in Neural Information Processing Systems*. 5254–5264.

[71] Sebastian Koos. 2012. What drives political consumption in Europe? A multilevel analysis on individual characteristics, opportunity structures and globalization. *Acta Sociologica* 55, 1 (March 2012), 37–57. https://doi.org/10.1177/0001699311431594

[72] Robert E Kraut, Paul Resnick, Sara Kiesler, Moira Burke, Yan Chen, Niki Kittur, Joseph Konstan, Yuqing Ren, and John Riedl. 2012. *Building successful online communities: Evidence-based social design.* Mit Press.

[73] Bogdan Kulynych, Rebekah Overdorf, Carmela Troncoso, and Seda Gürses. 2020. POTs: protective optimization technologies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 177–188.

[74] Shamanth Kumar, Reza Zafarani, and Huan Liu. 2011. Understanding User Migration Patterns in Social Media.. In *AAAI*, Vol. 11. 8–11.

[75] Shyong K Lam and John Riedl. 2004. Shilling recommender systems for fun and profit. In *Proceedings of the 13th international conference on World Wide Web*. 393–402.

[76] Cliff Lampe, Nicole B. Ellison, and Charles Steinfeld. 2008. Changes in Use and Perception of Facebook. In *Proceedings of the 2008 conference on Computer supported cooperative work*. 721–730.

[77] Cliff Lampe, Jessica Vitak, and Nicole Ellison. 2013. Users and nonusers: Interactions between levels of adoption and social capital. In *Proceedings of the 2013 conference on Computer supported cooperative work*. 809–820.

[78] Jong-Seok Lee and Dan Zhu. 2012. Shilling attack detection—a new approach for a trustworthy recommender system. *INFORMS Journal on Computing* 24, 1 (2012), 117–131. Publisher: INFORMS.

[79] Min Kyung Lee, Daniel Kusbit, Evan Metsky, and Laura Dabbish. 2015. Working with machines: The impact of algorithmic and data-driven management on human workers. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 1603–1612.

[80] Tuukka Lehtiniemi and Minna Ruckenstein. 2018. The social imaginaries of data activism. *Big Data & Society* 6, 1 (2018), 2053951718821146.

[81] Bo Li, Yining Wang, Aarti Singh, and Yevgeniy Vorobeychik. 2016. Data poisoning attacks on factorization-based collaborative filtering. In *Advances in neural information processing systems*. 1885–1893.

[82] Hanlin Li, Bodhi Alarcon, Sara M. Espinosa, and Brent Hecht. 2018. Out of Site: Empowering a New Approach to Online Boycotts. *Proceedings of the 2018 Computer-Supported Cooperative Work and Social Computing (CSCW'2018 / PACM)* (2018).

[83] Hanlin Li and Brent Hecht. 2020. 3 Stars on Yelp, 4 Stars on Google Maps: A Cross-Platform Examination of Restaurant Ratings. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW (2020).

[84] Hanlin Li, Nicholas Vincent, Janice Tsai, Jofish Kaye, and Brent Hecht. 2019. How do people change their technology use in protest?: Understanding "protest users". *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 87.

[85] Jiwei Li, Myle Ott, Claire Cardie, and Eduard Hovy. 2014. Towards a general rule for identifying deceptive opinion spam. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 1566–1576.

[86] Kim Lyons. 2021. FTC settles with photo storage app that pivoted to facial recognition. *The Verge* (Jan. 2021). https://www.theverge.com/2021/1/11/22225171/ftc-facial-recognition-ever-settled-paravision-privacy-photos Publisher: The Verge.

[87] Helen Margetts, Peter John, Scott Hale, and Taha Yasseri. 2015. *Political turbulence: How social media shape collective action.* Princeton University Press.

[88] Arunesh Mathur, Jessica Vitak, Arvind Narayanan, and Marshini Chetty. 2018. Characterizing the use of browser-based blocking extensions to prevent online tracking. In *Fourteenth Symposium on Usable Privacy and Security ({SOUPS} 2018)*. 103–116.

[89] J Nathan Matias. 2016. Going dark: Social factors in collective action against platform operators in the Reddit blackout. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 1138–1151.

[90] Connor McMahon, Isaac L Johnson, and Brent Hecht. 2017. The Substantial Interdependence of Wikipedia and Google: A Case Study on the Relationship Between Peer Production Communities and Information Technologies.. In *ICWSM*. 142–151.

[91] Stefania Milan and Lonneke Van der Velden. 2016. The alternative epistemologies of data activism. *Digital Culture & Society* 2, 2 (2016), 57–74.

[92] Bamshad Mobasher, Robin Burke, Runa Bhaumik, and Chad Williams. 2005. Effective attack models for shilling item-based collaborative filtering systems. In *Proceedings of the WebKDD Workshop*. Citeseer, 13–23.

[93] Yaroslav Nechaev, Francesco Corcoglioniti, and Claudio Giuliano. 2017. Concealing Interests of Passive Users in Social Media.. In *BlackMirror@ ISWC*.

[94] Edward Newell, David Jurgens, Haji Mohammad Saleem, Hardik Vala, Jad Sassine, Caitrin Armstrong, and Derek Ruths. 2016. User Migration in Online Social Networks: A Case Study on Reddit During a Period of Community Unrest.. In *ICWSM*. 279–288.

[95] Safiya Umoja Noble. 2018. *Algorithms of oppression: How search engines reinforce racism.* nyu Press.

[96] Myle Ott, Claire Cardie, and Jeff Hancock. 2012. Estimating the prevalence of deception in online review communities. In *Proceedings of the 21st international conference on World Wide Web*. 201–210.

[97] Kari Paul. 2020. Prime Day: activists protest against Amazon in cities across US. *the Guardian* (Apr 2020). https://www.theguardian.com/technology/2019/jul/15/prime-day-activists-plan-protests-in-us-cities-and-a-boycott-of-e-commerce-giant

[98] Business Paul R. La Monica. 2020. Tech's magnificent seven are worth $7.7 trillion. https://www.cnn.com/2020/08/20/investing/faang-microsoft-tesla/index.html [Online; accessed 6. Oct. 2020].

[99] Nikolaos Pitropakis, Emmanouil Panaousis, Thanassis Giannetsos, Eleftherios Anastasiadis, and George Loukas. 2019. A taxonomy and survey of attacks against machine learning. *Computer Science Review* 34 (Nov. 2019), 100199. https://doi.org/10.1016/j.cosrev.2019.100199 Publisher: Elsevier.

[100] Laura Portwood-Stacer. 2013. Media refusal and conspicuous non-consumption: The performative and political dimensions of Facebook abstention. *New Media & Society* 15, 7 (2013), 1041–1057.

[101] Eric A Posner and E Glen Weyl. 2018. *Radical Markets: Uprooting Capitalism and Democracy for a Just Society.* Princeton University Press.

[102] Barbara Prainsack. 2019. Data donation: How to resist the iLeviathan. In *The ethics of medical data donation*. Springer, Cham, 9–22.

[103] Alexander J Quinn and Benjamin B Bederson. 2011. Human computation: a survey and taxonomy of a growing field. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 1403–1412.

[104] Adam Satariano. 2020. What the G.D.P.R., Europe's Tough New Data Law, Means for You. *N.Y. Times* (May 2020). https://www.nytimes.com/2018/05/06/technology/gdpr-european-privacy-law.html Publisher: The New York Times Company.

[105] Christine Satchell and Paul Dourish. 2009. Beyond the user: use and non-use in HCI. In *Proceedings of the 21st Annual Conference of the Australian Computer-Human Interaction Special Interest Group: Design: Open 24/7*. 9–16.

[106] Devansh Saxena, Patrick Skeba, Shion Guha, and Eric PS Baumer. 2020. Methods for Generating Typologies of Non/use. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW1 (2020), 1–26.

[107] Sarita Yardi Schoenebeck. 2014. Giving up Twitter for Lent: how and why we take breaks from social media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 773–782.

[108] Roy Schwartz, Jesse Dodge, Noah A. Smith, and Oren Etzioni. 2020. Green AI. *Commun. ACM* 63, 12 (Nov. 2020), 54–63. https://doi.org/10.1145/3381831

[109] Neil Selwyn. 2003. Apart from technology: understanding people's non-use of information and communication technologies in everyday life. *Technology in society* 25, 1 (2003), 99–116.

[110] Alana Semuels. 2017. Why #DeleteUber and Other Boycotts Matter . *Atlantic* (Feb 2017). https://www.theatlantic.com/business/archive/2017/02/why-deleteuber-and-other-boycotts-matter/517416

[111] Rijurekha Sen, Sohaib Ahmad, Amreesh Phokeer, Zaid Ahmed Farooq, Ihsan Ayyub Qazi, David Choffnes, and Krishna P Gummadi. 2017. Inside the walled garden: Deconstructing facebook's free basics program. *ACM SIGCOMM Computer Communication Review* 47, 5 (2017), 12–24.

[112] Ali Shafahi, W Ronny Huang, Mahyar Najibi, Octavian Suciu, Christoph Studer, Tudor Dumitras, and Tom Goldstein. 2018. Poison frogs! targeted clean-label poisoning attacks on neural networks. In *Advances in Neural Information Processing Systems*. 6103–6113.

[113] Shawn Shan, Emily Wenger, Jiayun Zhang, Huiying Li, Haitao Zheng, and Ben Y. Zhao. 2020. Fawkes: Protecting Privacy against Unauthorized Deep Learning Models. In *29th {USENIX} Security Symposium ({USENIX} Security 20)*. 1589–1604.

[114] Márcio Silva, Lucas Santos de Oliveira, Athanasios Andreou, Pedro Olmo Vaz de Melo, Oana Goga, and Fabrício Benevenuto. 2020. Facebook Ads Monitor: An Independent Auditing System for Political Ads on Facebook. In *Proceedings of The Web Conference 2020*. 224–234.

[115] Jacob Steinhardt, Pang Wei W Koh, and Percy S Liang. 2017. Certified defenses for data poisoning attacks. In *Advances in neural information processing systems*. 3517–3529.

[116] Stefan Stieger, Christoph Burger, Manuel Bohn, and Martin Voracek. 2013. Who commits virtual identity suicide? Differences in privacy concerns, internet addiction, and personality between Facebook users and quitters. *Cyberpsychology, Behavior, and Social Networking* 16, 9 (2013), 629–634. Publisher: Mary Ann Liebert, Inc. 140 Huguenot Street, 3rd Floor New Rochelle, NY 10801 USA.

[117] Farnaz Tahmasebian, Li Xiong, Mani Sotoodeh, and Vaidy Sunderam. 2020. Crowdsourcing under data poisoning attacks: A comparative study. In *IFIP Annual Conference on Data and Applications Security and Privacy*. Springer, 310–332.

[118] Catalina L Toma and Jeffrey T Hancock. 2010. Reading between the lines: linguistic cues to deception in online dating profiles. In *Proceedings of the 2010 ACM conference on Computer supported cooperative work*. 5–8.

[119] Carmela Troncoso. 2019. Keynote Address: PETs, POTs, and Pitfalls: Rethinking the Protection of Users against Machine Learning. USENIX Association, Santa Clara, CA.

[120] Max Van Kleek, Dave Murray-Rust, Amy Guy, Kieron O'Hara, and Nigel Shadbolt. 2016. Computationally Mediated Pro-Social Deception. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. Association for Computing Machinery, New York, NY, USA, 552–563. https://doi.org/10.1145/2858036.2858060

[121] Nicholas Vincent and Brent Hecht. 2021. Can "Conscious Data Contribution" Help Users to Exert "Data Leverage" Against Technology Companies? *Proceedings of the ACM on Human-Computer Interaction* CSCW.

[122] Nicholas Vincent and Brent Hecht. 2021. A Deeper Investigation of the Importance of Wikipedia Links to the Success of Search Engines. *Proceedings of the ACM on Human-Computer Interaction* CSCW (2021).

[123] Nicholas Vincent, Brent Hecht, and Shilad Sen. 2019. "Data Strikes": Evaluating the Effectiveness of New Forms of Collective Action Against Technology Platforms. In *Proceedings of The Web Conference 2019*.

[124] Nicholas Vincent, Isaac Johnson, and Brent Hecht. 2018. Examining Wikipedia with a broader lens: Quantifying the value of Wikipedia's relationships with other large-scale online communities. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.

[125] Nicholas Vincent, Isaac Johnson, Patrick Sheehan, and Brent Hecht. 2019. Measuring the importance of user-generated content to search engines. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 13. 505–516.

[126] Nicholas Vincent, Yichun Li, Renee Zha, and Brent Hecht. 2019. Mapping the Potential and Pitfalls of "Data Dividends" as a Means of Sharing the Profits of Artificial Intelligence. *arXiv preprint arXiv:1912.00757* (2019).

[127] Kaveh Waddell. 2020. California's New Privacy Rights Are Tough to Use, Consumer Reports Study Finds. *Consum. Rep.* (Oct 2020). https://www.consumerreports.org/privacy/californias-new-privacy-rights-are-tough-to-use

[128] Daisuke Wakabayashi. 2018. California Passes Sweeping Law to Protect Online Privacy. *N.Y. Times* (Jun 2018). https://www.nytimes.com/2018/06/28/technology/california-online-privacy-law.html

[129] Hongyi Wen, Longqi Yang, Michael Sobolev, and Deborah Estrin. 2018. Exploring recommendations under user-controlled data filtering. In *Proceedings of the 12th ACM Conference on Recommender Systems*. 72–76.

[130] John Wilmhoff. 2017. Tom Brady literally owns the Jets, says Google search. https://www.espn.com/sportsnation/story/_/page/170727QTP_BradyOwnsJets/google-glitch-causes-tom-brady-appear-new-york-jets-owner%7D Publication Title: ESPN.

[131] David C Wilson and Carlos E Seminario. 2013. When power users attack: assessing impacts in collaborative recommender systems. In *Proceedings of the 7th ACM conference on Recommender systems*. 427–430.

[132] David C Wilson and Carlos E Seminario. 2014. Evil twins: Modeling power users in attacks on recommender systems. In *International Conference on User Modeling, Adaptation, and Personalization*. Springer, 231–242.

[133] Sally ME Wyatt. 2003. Non-users also matter: The construction of users and non-users of the Internet. *Now users matter: The co-construction of users and technology* (2003), 67–79.

[134] Sean Xin Xu and Xiaoquan (Michael) Zhang. 2013. Impact of Wikipedia on Market Information Environment: Evidence on Management Disclosure and Investor Reaction. *MIS Quarterly* 37, 4 (Dec 2013), 1043–1068. http://www.jstor.org/stable/43825781

[135] Haoqi Zhang, Andrés Monroy-Hernández, Aaron Shaw, Sean A Munson, Elizabeth Gerber, Benjamin Mako Hill, Peter Kinnaird, Shelly D Farnham, and Patrick Minder. 2014. WeDo: end-to-end computer supported collective action. In *Eighth International AAAI Conference on Weblogs and Social Media*.

[136] Renjie Zhou, Samamon Khemmarat, and Lixin Gao. 2010. The impact of YouTube recommendation system on video views. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*. ACM, 404–410.