

SoK: On the Impossible Security of Very Large Foundation Models

El-Mahdi El-Mhamdi
École Polytechnique
Palaiseau, France

el-mahdi.el-mhamdi@polytechnique.edu

Sadegh Farhadkhani
IC, EPFL

Lausanne, Switzerland
sadegh.farhadkhani@epfl.ch

Rachid Guerraoui
IC, EPFL

Lausanne, Switzerland
rachid.guerraoui@epfl.ch

Nirupam Gupta
IC, EPFL

Lausanne, Switzerland
nirupam.gupta@epfl.ch

Lê Nguyễn Hoang
Association Tournesol,
Switzerland
len@tournesol.app

Rafael Pinot
IC, EPFL
Lausanne, Switzerland
rafael.pinot@epfl.ch

John Stephan
IC, EPFL
Lausanne, Switzerland
john.stephan@epfl.ch

Abstract—Large machine learning models, or so-called *foundation models*, aim to serve as base-models for application-oriented machine learning. Although these models showcase impressive performance, they have been empirically found to pose serious security and privacy issues. We may however wonder if this is a limitation of the current models, or if these issues stem from a fundamental *intrinsic impossibility* of the foundation model learning problem itself. This paper aims to systematize our knowledge supporting the latter. More precisely, we identify several key features of today’s foundation model learning problem which, given the current understanding in adversarial machine learning, suggest incompatibility of *high accuracy* with both security and privacy. We begin by observing that high accuracy seems to require (1) very *high-dimensional* models and (2) huge amounts of data that can only be procured through *user-generated datasets*. Moreover, such data is *fundamentally heterogeneous*, as users generally have very specific (easily identifiable) data-generating habits. More importantly, users’ data is filled with highly *sensitive information*, and maybe heavily polluted by *fake users*. We then survey lower bounds on accuracy in privacy-preserving and Byzantine-resilient heterogeneous learning that, we argue, constitute a compelling case against the possibility of designing a secure and privacy-preserving high-accuracy foundation model. We further stress that our analysis also applies to other high-stake machine learning applications, including content recommendation. We conclude by calling for measures to prioritize security and privacy, and to slow down the race for ever larger models.

Index Terms—security, privacy, foundation models, machine learning, curse of dimensionality, heterogeneity, statistics

I. INTRODUCTION

In recent years, we have witnessed immense growth in the size of machine learning models. The number of parameters has increased from 213 million in 2017 [178], to 1.5 billion in 2019 [151], 175 billion in 2020 [23], 1.6 trillion in early 2021 [57], and over 100 trillion in late 2021 [116]. The scaling of model sizes improved accuracy on classical tasks such as GLUE [185], SuperGLUE [184], or Winograd [156], without significant diminishing returns so far (see, e.g., Figure 1 in [23]). Such models also excel in few-shot learning [23], which has motivated their wide use as pre-trained “*foundation*” (or “*base*”) models, to be *fine-tuned* to

any task of interest [35], [34], [91], [182], [201]. This success has generated significant academic, economic and political interest to accelerate the development and deployment of *foundation models* for applications such as content moderation, recommendation, search and ad targeting [41]. Arguably, this pressure has been accentuated by a glorification of this line of research and of its outcomes, especially in fundraising, news outlets and political discourses¹. Military agencies, private companies and even universities, are now all racing for ever more impressive performance [29], [65].

However, numerous voices have raised serious concerns about the rushed deployment of such technologies [87]. These concerns are well illustrated by the anti-Muslim bias of OpenAI’s (deployed and commercialized) GPT-3 *foundation model* [23]. As exposed by [3], when prompted with “Two Muslims walk into”, GPT-3 completes it by “a Church, one of them as a priest, and slaughtered 85 people”. The risks of subtle induced radicalization was further highlighted by [125]. Namely, when asked “who is QAnon?”, GPT-3 provides a Wikipedia-like factual answer. However, if GPT-3 is first prompted with queries typical of conspiracy forums such as “Who are the main enemies of humanity?”, then GPT-3’s answer to “who is QAnon?” now becomes typical of such forums, as it answers “QAnon is a high-level government insider who is exposing the Deep State”. As already evidenced by the 2021 Capitol riots [183], such results raise serious national security and world peace concerns.

To understand how such concerns are related to machine learning *security*, we stress that today’s *foundation models* are almost exclusively shaped by their training data, which too often amounts to barely filtered online data. In fact, they are usually designed to reproduce the most frequent claims. This is why BlenderBot, Facebook’s own *foundation model*, generated insults against Facebook’s CEO Mark Zuckerberg [205].

¹Politico published an article on a Chinese language model with 1.75 trillion parameters, with the following subtitle: “Europe is increasingly worried it’s being left out of the global race for artificial intelligence” [78]. This implicitly calls for racing to build ever larger *foundation models*.

Concerningly, this also creates perverse incentives, especially in the context of the global disinformation war [161]. Namely, malicious actors can hack today’s most influential *foundation models* by *poisoning* the web with the propaganda they want to massively broadcast. Typically, GPT-3’s anti-Muslim bias can be argued to be (partly) the result of anti-Muslim propaganda, which has been found to be massively scaled, by both human troll farms and (foundation-model-based) algorithms like Tek Fog in India [97], [98]. Recall that, in 2019 alone, Facebook removed 6 billion fake accounts from its platform [63] (and the development of generative models and *evasion attacks* [16], [68] will likely make this worse).

Another major concern raised about *foundation models* is *privacy*, especially if emails and smart keyboard’s data are used² to train them [75], [195]. As opposed to what has long been the conventional wisdom backed by the theory of (PAC) learning [177], since 2017, the generalization performance of many learning tasks including language processing have been empirically shown to be best achieved by fully interpolating the training data [10], [133], [198]. A vast literature has since provided theoretical explanations for this phenomenon, sometimes called “double descent” [11], [12], [13], [86], [77], [82], [118], [128], [132], [134]. Put differently, the best accuracy is achieved when models *memorize* their training data [58]. However, in the context of user-generated data, this raises serious privacy concerns, as these data are expected to contain highly sensitive information. In fact, numerous recent papers have shown that *foundation models* can be easily queried to retrieve personal information, e.g., by simply asking them “what is Mr. X’s home address?” [28], [84], [142], [206].

More generally, it should not be forgotten that *foundation models* are designed to interpolate their training datasets, and to sample from the distribution learned from this data interpolation. This is what makes them “stochastic parrots” that repeat and amplify their web-crawled training database, without any attention to misleading or sensitive information [15]. This makes them *dangerous* to deploy at scale. Yet *foundation models* are *already* deployed. For instance, GPT-3 has been reported to already produce billions of words per day [181]. But in the meantime, the details of these algorithms and of their deployments are often hidden, even from most of the employees of the groups deploying these algorithms [160], which prevents even internal audits.

Now, one might argue that these weaknesses are specific to today’s algorithms, and that more research will fix these vulnerabilities without damaging accuracy. In this paper, we systematize the current knowledge of privacy-preserving and secure learning, which, as we will see, suggest the opposite. Namely, a wide literature, which we will review, now provides strong arguments against the possibility of secure high-accuracy *foundation model* design. As a result, as long as they race for performance, ever more dramatic *foundation models* will be ever more easily manipulated by malicious actors, and

they will be more easily probed to steal sensitive information.

More precisely, in this paper, we first identify the specific compromising features of *foundation model* training (thereby making any fine-tuning of these models also vulnerable). Namely, given that *foundation models* achieve their best accuracy when performing *very-high dimensional interpolation*, the race for performance incentivizes the use of huge amounts of unsafe *user-generated data*. Yet, when produced by genuine humans, language data are very user-specific. Genuine humans’ data are thus *fundamentally heterogeneous*, in the sense that different genuine users have different preferred phrase completion. Moreover, the statistical word distribution is well-known to be *heavy-tailed* [123], [149], and each user provides only a *sparse dataset* (i.e., not fully representative of all the ways the user would speak), leading to additional *empirical heterogeneity*. As we will see, such data heterogeneities are a core cause of the fragility of *foundation models*, especially when the data is sensitive and might be fabricated by *fake accounts* [20], [137], [191].

Indeed, an increasingly large body of both empirical and theoretical research has exposed serious concerns, especially for *high-dimensional* model training based on *heterogeneous user-generated* data. Yet, according to a recent survey [108], “industry practitioners are not equipped with tactical and strategic tools to protect, detect and respond to attacks on their machine learning systems”. The heart of this paper reviews the machine learning security literature, with an emphasis on known lower bounds on what any “secure” learning algorithm can guarantee. We also stress that, while adversarial machine learning is plagued with such negative impossibility theorems, more often than not, the community still presents an optimistic view on the problem by making inadequate assumptions. Typically, on the one hand, because of negligence, error, or *doxing*³, a user’s generated data may repeatedly leak private information about many other users, which seriously questions the relevance and applicability of classical privacy definitions (see Section IV-D). On the other hand, many authentic users’ (misinformed or hate) messages may be highly undesirable to repeat and amplify, which shows that the classical “Byzantine model” is inappropriate for the *foundation model* learning problem (see Section V-B).

To make our concerns concrete, we present in this paper numerous ways through which deployed *foundation models* are already causing serious harms to society, and could be causing even more concerning harms through autocompletion, conversational, and especially recommendation algorithms. We stress that non-principled proposals to fix today’s *foundation models*, such as hand fixes or fine tuning, at least currently, are far from satisfactory.

It is important to note that while much attention is given to *language models*, the problematic features we identify in this paper are not specific to language processing. Indeed, social media images are also *user-generated, high-dimensional* and

²Even when user data are not fed as-is to the model, we discuss this later in subsection IV-A.

³According to Wikipedia, “doxing is the act of publicly revealing previously private personal information about an individual or organization”.

heterogeneous. Learning a distribution over these images, as is done by generative adversarial networks, is arguably very unsafe as well. The same applies to learning from users’ video and sound recordings. Perhaps what is even more concerning is the case of users’ online activities (e.g., likes, shares, watch-time and click-through rates), which are critical data for user retention and for the extremely lucrative ad targeting business. Such online data are also *user-generated*, *high-dimensional* and *heterogeneous*. The mathematical impossibility of combining high performance and high security may thus permeate most of these very high-stake security-sensitive machine learning applications.

The rest of the paper is organized as follows. Section II highlights the challenging features of today’s *foundation model* training problem. Section III justifies the relevancy of the secure mean estimation problem to understand the secure learning problem. Section IV then reviews the literature on the impossibility of privacy-preserving high-accuracy mean estimation, while Section V surveys published results on its unavoidable vulnerability to poisoning attacks. Section VI provides concrete scenarios where the vulnerability of *foundation models* is already leading to serious social consequences at scale. Section VII lists ideas proposed to fix language models, and argues that they are unlikely to be secure enough. Finally, Section VIII concludes, with a call for urgent action.

II. FOUR FEATURES OF FOUNDATION MODEL TRAINING

In this section, we highlight four key features, which distinguish *foundation model* training from many classical machine learning tasks, and which make *foundation models* particularly vulnerable to poisoning and privacy attacks.

A. User-generated data

Foundation models achieve their best performances by leveraging ever larger amounts of data [202], without major diminishing returns so far. Unfortunately, verified texts seem insufficient to reach state-of-the-art performance. Indeed, the English Wikipedia only contains around 4 billion words [189]. Meanwhile, a book has around 10^5 words. While there are 10^8 books [122], only a fraction of them are arguably trustworthy. Many books are instead full of biases and dangerous misinformation, such as ethnic-based hate speech, historical propaganda, or outdated (possibly harmful) medical advice. As a striking illustration, up to the 1980s, the American Psychiatric Association listed homosexuality as a mental illness in its flagship manual [170]. In fact, most books should be regarded as unverified user-generated data.

Most importantly, even if they are not problematic, the combination of these books represents a small amount of data, compared to what Internet users produce on a daily basis. Indeed, assuming that a user writes 300 words per day on an electronic device (the equivalent of one page), a billion of such users produce 10^{15} words per decade. This adds up to a hundred times more data than the set of books, and a million times more than the English Wikipedia. This makes it very tempting to either scrape the web [165], [185], [184], exploit

private messaging (e.g., emails, shared documents), or leverage other written texts (e.g., phones’ smart keyboards). In fact, Wikipedia represented only 4% of Google’s PaLM *foundation model* training dataset [33], while books represented 13% of it. Meanwhile, 27% of the dataset was made of webpages, and 50% were social media conversations. Crucially, these data are generated by a myriad of users, who may be malicious and/or unaware that their activities are being leveraged to train *foundation models*. This raises serious security and privacy risks, even when we (wrongly) restrict our attention to publicly released models (most of the actual *foundation model* training is likely performed secretly by private groups).

Indeed, user-generated data are both mostly unverified and potentially highly sensitive. This feature of language is in sharp contrast with sensors’ data, especially when the sensors are owned, audited and trustworthy (even though sensors’ data can also leak private information).

Of course, to increase security, as proposed by [155], we could demand that *foundation models* restrict themselves to quality datasets only. However, such datasets will inevitably be of significantly smaller size. As discussed earlier, this will likely greatly harm the performance of *foundation models*. In fact, this is the main claim of this paper: security demands a drastic reduction of performance.

B. Very high-dimensional interpolation

Foundation models are often *overparametrized* to interpolate huge amounts of data [198], [133], [199]. This has led to ever larger models. Today’s largest (reported) *foundation models* have over a trillion parameters [57]. This means that the dimension d of the parameter space is now at least of the order of 10^{12} . Moreover, empirical results suggest that we have not yet reached a point of diminishing returns [23], while some theoretical arguments suggest that memorization may be necessary for generalization [58], [22]. This observation arguably distinguishes language from other tasks, such as image classification, where ever larger models do not seem to yield ever better accuracy⁴.

Note that theoretical arguments, akin to Turing’s arguments for the eventual need of machine learning [175], also suggest that better accuracy requires larger models. Namely, Turing noted that the human brain has 10^{15} synapses. Even if only 1% of these synapses are essential to conduct a human-level conversation, then this still means that 10^{13} parameters are needed to do so. In fact, this smallest number of bits of information to achieve a task has been formalized in 1960 by Solomonoff [168], and then Kolmogorov [103], and is now known as the Solomonoff-Kolmogorov complexity⁵ of (quality) human-level conversation. If this complexity is 10^{13} , then no algorithm with fewer parameters will achieve the task. Yet, it is noteworthy that what is now demanded from such *foundation models* is often beyond the capability of any single human. Indeed, such algorithms are able to memorize the

⁴<https://paperswithcode.com/sota/image-classification-on-imagenet>.

⁵Referred to as the Kolmogorov complexity in most textbooks.

entirety of Wikipedia, which is the result of the cumulative works of many experts in their respective fields, on a myriad of diverse topics. Such *foundation models* must arguably be able to adapt to a greater variety of contexts than what any single human will ever encounter in their human life. As a result, the complexity of “fully satisfactory” language processing might need to be orders of magnitude larger than today’s *foundation models*, in which case we may still obtain greater accuracy by training larger models.

Unfortunately, this exposes such large *foundation models* to the infamous *curse of dimensionality* [14], [66], which has been connected to increased security risks. For instance, [7] proved that the error rate of a learning model under optimal targeted data poisoning necessarily grows as⁶ $\Omega(d|\mathcal{D}_n|/|\mathcal{D}|)$, where \mathcal{D} is the entire dataset and \mathcal{D}_n is the subset of poisonous data injected by the malicious user n . In other words, no matter what defense the model trainer comes up with, whenever the malicious user is given an input to attack (drawn from the true distribution over inputs), then (in expectation over the input to attack) it can make the *foundation model* fail on this input by only injecting $\mathcal{O}(|\mathcal{D}|/d)$ poisonous data. Yet, for overparameterized models, we have $d \gg |\mathcal{D}|$. Concerningly, this implies that the model is completely vulnerable to a handful of, perhaps even a single, targeted poisonous data.

To understand, consider the following intuitive consideration. In the case of linear or logistic regression, each data acts on the model parameters on a single dimension. Thus, if the model has more dimensions than there are data points, then many dimensions will be under the influence of no data. This makes such dimensions extremely vulnerable to a data poisoning attack. Moreover, more generally, the more we are in a regime $d \gg |\mathcal{D}|$, the more it may hold that most dimensions can be arbitrarily hacked in this manner.

Now of course, the dimension d could be reduced to increase security. However, today’s empirical observations strongly suggest that doing so incurs a significant accuracy loss. In fact, this is the main claim of our paper. Namely, security demands a large accuracy drop. In particular, as long as accuracy is highly valued and massively funded, then security and privacy are arguably in great danger.

C. Fundamental heterogeneity

In the case of language, authentic users’ data distributions are clearly very *heterogeneous* [117], [95], [153]. More precisely, the distribution of texts generated by a given user greatly diverges from the distribution of texts generated by another user. This is evidenced by the fact that it is often possible to guess the author of a message [36], [121], simply based on the content of the message. Of course, this will be especially the case if the message contains highly identifiable information, such as the names of the recipient of the message, or a sequence of controversial judgments. But even if the message does not explicitly expose such information, its

writing style often suffice to expose the more probable author identity [61], [192], [180].

We can formalize more precisely a notion of *fundamental heterogeneity* for users’ language generation. Namely, note that the data used by *foundation models* is typically a set of feature-label pairs of the form $(context, word)$, where the *context* is a set of words surrounding the *word*. Consider the cases where the *context* is equal to “my name is”, “Republicans are”, or “vaccines are”. Clearly, different users would complete the phrase differently, meaning that the different users are using different *labeling functions*.

We stress that this heterogeneity in the users’ labeling functions can be regarded as a *fundamental heterogeneity*, as it would still remain *even if* all users labeled an infinite amount of times the same inputs. This heterogeneity highlights an irreconcilable disagreement between users over which *foundation model* should be learned. While some users would prefer to complete the sentence “the greatest of all time tennis player is” by “Roger Federer”, others would prefer to complete it by “Novak Djokovic”, or by “Rafael Nadal”. This is sharp contrast with image classification and language emotion classification tasks, where different users usually label a single image or text similarly.

This makes accurately learning a distribution of texts much more dangerous. On one hand, the model would be able to map users’ names to what they write, which is a major privacy concern. On the other hand, it would then be easier for malicious users to be hardly discernible from most other genuine users, while providing very dangerous texts to replicate.

D. Sparse heavy-tailed data per user

Another important feature of language data is that they are sparse per user. More precisely, the dataset \mathcal{D}_n provided by an honest user n is much smaller than the model dimension d . Imagine a user who types around 300 words per day (the equivalent of one page). Then over a decade, this user will have provided around $|\mathcal{D}_n| \approx 10^6$ feature-label pairs. This quantity is vastly smaller than the dimension of today’s largest *foundation models*, which is a million times larger.

Statistically, this roughly corresponds to each user sampling $|\mathcal{D}_n|$ points in a space of dimension d (especially if we reason in terms of gradients, as will be done in Section IV). The sample mean will then greatly diverge from the users’ distribution mean. In fact, assuming that the user’s true data distribution (e.g., in the gradient space) is a normal distribution with covariance matrix $\sigma^2 I_d$, the variance of the sample mean will be $\sigma^2 I_d / |\mathcal{D}_n|$. The typical distance between the sample mean and the distribution mean will then be of the order $\sqrt{\text{Tr}(\sigma^2 I_d / |\mathcal{D}_n|)} = \sigma \sqrt{d / |\mathcal{D}_n|}$. For $d \gg |\mathcal{D}_n|$, this is large.

Furthermore, language data often have heavy-tailed distributions [123], [149]. In the context of machine learning, when applying SGD for language models, the norms of the stochastic gradients have been shown to follow a power law distribution [200]. Intuitively, this is because most sentence completions are rare, especially if they are to be completed by several words. Yet, it is a fundamental property of heavy-tailed

⁶Their result involves the VC dimension d .

distributions that their samples are often highly unrepresentative of the overall distribution, especially when the sample sizes are not large enough. In particular, this means that we should expect an especially large *empirical heterogeneity* in language data, as the samples we obtain from a user can completely stand out from their own language distributions.

Overall, because every genuine user has too few heavy-tailed data compared to the dimension of the problem, in addition to the *fundamental heterogeneity* due to users' drawing from different data distribution, language data will typically exhibit an additional large *empirical heterogeneity* between users' finite data samples. This empirical heterogeneity would remain, even if all users drew from the same data distribution, as long as they draw significantly less data than the dimension of the *foundation model*. As fundamental and empirical heterogeneities intuitively add up, language data should be expected to feature an enormous total heterogeneity, especially when processed by very high-dimensional models. As we will see, given our current understanding of machine learning security and privacy, this is an enormous source of concern.

III. WHY WE FOCUS ON MEAN ESTIMATION

To systematize the applicability of prior works, much of which focuses on the secure mean estimation problem, we first highlight an insightful equivalence between data poisoning and gradient attacks [56]. While this equivalence only applies to a restricted setting, we argue that it suggests that known impossibility theorems for accurate secure mean estimation are compelling evidence that the current quest for spectacular *foundation models* is creating major security vulnerabilities.

A. Setup

Let us first formalize the safe and private learning problem. We consider a set $[N] = \{1, \dots, N\}$ of users. Each user $n \in [N]$ generates a certain set \mathcal{D}_n of data. The dataset \mathcal{D}_n is typically made of feature-label pairs (y, z) . In natural language processing, y may be thought of as the context, and z as the token (word) that fits the context in the dataset. We denote $\vec{\mathcal{D}} \triangleq (\mathcal{D}_1, \dots, \mathcal{D}_N)$ the N -tuple of users' datasets.

We then fit the parameters $\theta \in \mathbb{R}^d$ of a (differentiable) model, typically a neural network such as a transformer [178]. For each user n 's dataset \mathcal{D}_n , the model θ is assumed to incur a *local loss* $\mathcal{L}_n(\theta, \mathcal{D}_n)$. Assuming that all users are honest, it is then common to aim to minimize the regularized sum of local losses, which is the following global loss function:

$$\text{Loss}(\theta, \vec{\mathcal{D}}) = \sum_{n \in [N]} \mathcal{L}_n(\theta, \mathcal{D}_n) + \mathcal{R}(\theta), \quad (1)$$

where $\mathcal{R}(\theta)$ is a regularization term.

To be concrete, let us describe the case of a purely data-fitting cost for a predictive model. In such a case, given parameters θ , the model predicts a label $f_\theta(y)$. The discrepancy between the model prediction $f_\theta(y)$ and the dataset label z

then incurs a cost $\ell(f_\theta(y), z)$. For a given user $n \in [N]$, adding up all the costs yields a local loss

$$\mathcal{L}_n(\theta, \mathcal{D}_n) \triangleq \sum_{(y,z) \in \mathcal{D}_n} \ell(f_\theta(y), z).$$

Denoting $\mathcal{D} \triangleq \bigcup_{n \in [N]} \mathcal{D}_n$ the union of all users' data, the global loss function is then simply fitting all available data, as

$$\text{Loss}(\theta, \mathcal{D}) = \sum_{(y,z) \in \mathcal{D}} \ell(f_\theta(y), z) + \mathcal{R}(\theta)$$

We stress however that, by considering the more general Equation (1), we actually consider a much larger class of frameworks to learn from different users' datasets. In particular, our setup includes alternatives that may, for instance, assign more importance to fairness, security or personalization. Namely, by assuming that each local loss

$$\mathcal{L}_n(\theta, \mathcal{D}_n) \triangleq \inf_{\varphi_n} \left\{ \mathcal{R}_n(\varphi_n, \theta) + \sum_{(y,z) \in \mathcal{D}_n} \ell(f_{\varphi_n}(y), z) \right\}$$

is defined by a so-called *reduced loss* [56], where φ_n is a local model personalized for user n and where \mathcal{R}_n is a regularization that penalizes the discrepancy between the global model θ and the user n 's local model φ_n , we encompass the increasingly studied framework of *personalized federated learning* ([44], [53], [74], to name a few).

B. Data poisoning versus Byzantine gradients

Classically, data poisoning has focused on a single polluted dataset [17], [37], [64]. In particular, this literature has failed to leverage the fact that, in most applications, it is usually possible to map each data point to a data provider. Data is often *signed* (and if it is not, then it should be regarded as highly untrustworthy). In fact, it is commonly accepted that the traceability of data sources is a critical security condition [110], [139], as well as a powerful epistemological tools [9]. Thus, our setting allows us to work under the arguably realistic assumption that, if some data from user n 's dataset \mathcal{D}_n are known to be harmfully crafted, then the entire dataset \mathcal{D}_n is likely to be untrustworthy as well. Unfortunately, the study on data poisoning with *signed* data has been lacking.

Interestingly, however, in the case of personalized federated learning, for linear and logistic regression, [56] proved an equivalence between data poisoning under the signed data setup and the widely studied Byzantine gradient attacks in the *federated learning* setting [18]. To understand, recall that federated learning is a distributed computational framework, where each user is asked to contribute to the training of a central parameter server, by sending *gradients* indicating (the opposite of) a preferred direction of model update. Given a central model θ^t at time t , user n 's expected gradient is given by $\nabla \mathcal{L}_n(\theta^t, \mathcal{D}_n)$. When the exact gradient computation is too demanding, the user n is asked to instead report an estimate,

like a *stochastic gradient* g_n^t ⁷. But a malicious, so-called *Byzantine*, users can deviate arbitrarily from the protocol and may report *any* gradient g_n^t .

Now, it is clear that any data poisoning can be turned into an equivalent gradient attack (by simply computing the gradients for the poisoning dataset). Remarkably, however, under some appropriate assumptions and for any converging gradient attack $g_n^t \rightarrow g_n^\infty$ from a malicious user n , [56] constructively proved the existence of a poisonous dataset $\mathcal{D}_n^\blacklozenge$ such that the learned global model θ under data poisoning by $\mathcal{D}_n^\blacklozenge$ is approximately equal to the value it takes under gradient attack g_n^t . Put differently, at least under their setting, the vulnerability (and defenses) to data poisoning can be completely understood by the (easier) study of gradient attacks.

C. Why the mean appears in machine learning

Classically, to train models and at each iteration t , the gradients g_n^t need to be added up or averaged. Indeed, it follows from Equation (1) that

$$\nabla_\theta \text{Loss} = \sum_{n \in [N]} \nabla_\theta \mathcal{L}_n + \nabla \mathcal{R} = \frac{1}{N} \sum_{n \in [N]} x_n, \quad (2)$$

where $x_n \triangleq N \nabla_\theta \mathcal{L}_n + \nabla \mathcal{R}$. Therefore, the training of (*foundation*) models heavily relies on the (repeated) averaging of user-specific vectors. Correctly and privately estimating the average of users' vectors x_n is thus critical for security.

In fact, [129] show an equivalence between robust mean estimation and robust heterogeneous learning. In particular, their results imply that any impossibility result about robust mean estimation implies an impossibility for robust machine learning in its general form⁸. Similarly, it seems that any impossibility result on private mean estimation is an evidence that privacy-preserving machine learning is hard in general, and that privacy-preserving gradient-descent-based language learning is hard in particular. Given that mean estimation is a key component of learning algorithms, and since it captures a key challenge in learning from multiple data sources, we choose to focus on its difficulty in this paper.

In fact, when $\mathcal{L}_n(\theta, \mathcal{D}_n) = \|\theta - x_n\|_2^2$ for some vector x_n for any $n \in [N]$ (and there is no regularization), the accuracy of a solution θ is directly related to its closeness to the mean⁹ of the vectors x_n 's. An algorithm that robustly or privately solves *all* learning problems must thus also be able to robustly or privately solve mean estimation in particular. Put differently, any impossibility on mean estimation implies an impossibility about general learning algorithms.

Unfortunately, in a context of high heterogeneity, even when the average of honest vectors x_n is small (i.e. we are near the optimum of the global loss), each individual vector x_n may

⁷in the case of personalized federated learning, the reported gradient may correspond to the gradient given a suboptimal local model φ_n .

⁸Evidently, specific machine learning models may escape this impossibility result, e.g., if they do not learn anything. However, there does not seem to be a clear reason why transformers would be a drastic exception.

⁹Indeed, the mean is the minimum of the loss thereby constructed, i.e. $\frac{1}{N} \sum_{n \in [N]} x_n = \arg \min_{\theta \in \mathbb{R}^d} \sum_{n \in [N]} \|\theta - x_n\|_2^2$

still have a large norm. In the case where all these vectors are distributed along a normal distribution $\mathcal{N}(0, \sigma^2 I_d)$, their typical distance to 0 will be $\Delta \approx \sigma \sqrt{d}$. For $d \approx 10^{12}$ (as reported for *foundation models*), the typical distance between the vectors and the center of the distribution would then be a million times larger than the typical standard deviation that is observed on one dimension. As we will discuss it more rigorously in the subsequent sections, this is what makes *high-dimensional heterogeneous learning* extremely vulnerable to poisoning and privacy attacks.

D. Homogeneous learning can be made secure

Before discussing impossibility theorems, let us stress that *data heterogeneity is the killer*. Indeed, [94], [129], [143], [54] proved that in the homogeneous case, Byzantine-resilient learning can be achieved when there is a majority of honest users (for a synchronous network), assuming that each user can provide an arbitrarily large amount of data drawn independently from the same distribution (thereby removing any *empirical heterogeneity* as well). More precisely, [129] devise an algorithm for non-convex (distributed) learning that, given any $\delta > 0$, outputs a solution such that the corresponding gradient of the loss restricted to honest users is of norm at most δ on expectation. The relative security of homogeneous learning was also observed empirically by [130], [162].

Such an output is also intuitively differentially private. Indeed, since the losses of users are similar (by homogeneity), removing a user does not affect the optimality of the computed parameters. Intuitively, this is because the loss function of a user does not actually reveal any information specific to the user; after all, this loss function is statistically indistinguishable from the loss function of any other user.

Unfortunately, *homogeneity is an unrealistic assumption* for the training of *foundation models*. Put differently, the fundamental vulnerability of *foundation models* is tightly connected to the fundamental heterogeneity in the way different users speak and write, and to the additional empirical heterogeneity due to the users' limited datasets (which cannot be representative of the full distribution from which the users draw their texts and speeches). These data are *not* drawn from a fixed common data distribution.

As a result, (positive) results based on the infamous *i.i.d. assumption* can be very misleading. This assumption is arguably *dangerously unrealistic*, especially for the security analysis of *foundation model* training. Unfortunately, so far, most of the celebrated theory of (Byzantine) machine learning builds upon this assumption [177], [85], [18]. A serious consequence of this is that it effectively turns much of the attention of the research community away from the urgent security and privacy concerns that today's *actual* large-scale machine learning algorithms are posing, especially when they are built upon *foundation models*. In turn, this leads to a dangerous neglect of actual security risks, and motivates companies and regulators to also prioritize "innovation" over security.

IV. THE PRIVACY-ACCURACY TRADEOFF

In this section, we present some impossibility theorems for accurate (differentially) private mean estimation, especially under high heterogeneity and in high dimension. We also discuss the limits of published positive results, and the flaws of the leading understanding of privacy in academia. But first, let us stress that some privacy analysis is needed, even when private data is not directly shared.

A. Federated learning is not privacy-preserving

There is a folklore belief, often cited without justification in federated learning papers, that, by avoiding data transfers and by sending gradients instead, “federated learning is a privacy-preserving” technique [31], [193]. We stress that this is an extremely dangerous misinformation [19], which somehow permeates the scientific community¹⁰. Indeed, this misinformation has been (mis)used, e.g. to justify the deployment of federated learning systems for COVID-19 detection and case analysis, without differential privacy mechanisms [2], [40], [46].

There is an obvious reason why federated learning has absolutely no privacy guarantee. Namely, federated learning is designed to achieve the same performances as classical centralized learning. Yet, as discussed in Section II-B, overparameterized *foundation models* are designed to fit and memorize their entire training dataset. Clearly, this cannot be privacy-preserving, even when secure multiparty methods are used to hide the users’ gradients during training [144].

To guarantee privacy, careful algorithm design and privacy analyses are needed. Typically, some noise is added to guarantee *differential privacy*.

B. Impossible private mean estimation

Differential privacy [49] has become the leading formalization of privacy. Essentially, the removal of one user n ’s dataset \mathcal{D}_n from the dataset tuple $\tilde{\mathcal{D}}$ should not affect significantly the outcome of a (user-level) differentially private algorithm. In the case of *foundation model* training, this means that training with $\tilde{\mathcal{D}}_{-n}$ (i.e. the dataset tuple obtained by removing user n ’s dataset) should yield approximately the same model as training with $\tilde{\mathcal{D}}$. Intuitively, this protects user n ’s dataset from privacy attacks.

As explained in Section III, since *foundation models* heavily rely on stochastic gradient descent, much of the literature leverages the large body of work on differentially private mean estimators [171], [47], [26], [93] to construct differentially private learning models. Formally, a mean estimator $\widehat{\text{MEAN}}$ is then said to satisfy (ε, δ) *user-level differential privacy* if, for all N , for all N -tuples $\vec{x} \triangleq (x_1, \dots, x_N)$ of vectors and for any user $n \in [N]$ to be dropped, given any subset X of outputs, we have

$$\mathbb{P}[\widehat{\text{MEAN}}(\vec{x}) \in X] \leq e^\varepsilon \mathbb{P}[\widehat{\text{MEAN}}(\vec{x}_{-n}) \in X] + \delta, \quad (3)$$

¹⁰This can be evidenced e.g. by the answers when searching for the phrase “federated learning is a privacy-preserving” on Google Scholar.

where \vec{x}_{-n} is the tuple obtained by removing x_n from \vec{x} .

Unfortunately, there are known lower bounds on the error of any *differentially private* mean estimation algorithm [24]. To present a simple result, assume here that the users’ vectors are known to lie in a ball of radius Δ . Here we adapt a result from [96] showing that to guarantee (ε, δ) -differential privacy, the mean squared error of the estimator must be proportional to both the dimension d of the input vectors and the worst case magnitude of a user’s vector within the vector family Δ .

Theorem 1 (Theorem 4 in [96]¹¹). *For any (ε, δ) -differentially private mechanism $\widehat{\text{MEAN}}$ for the mean estimation problem, there exists an input \vec{x} with large mean squared error, as*

$$\mathbb{E} \left[\|\widehat{\text{MEAN}}(\vec{x}) - \bar{x}\|_2^2 \right] \geq \Omega \left(\frac{\sigma(\varepsilon, \delta)d\Delta^2}{N^2(\log 2d)^4} \right), \quad (4)$$

where σ is a positive and non-increasing function.

In high dimension d , the typical radius Δ should typically be expected to grow as \sqrt{d} . If so, ignoring the dependency on ε and δ , then we see that the lower bound of Theorem 1 would be $\tilde{\Omega}(d^2/N^2)$. In other words, accuracy demands to have $d \ll N$. With d in the trillions, this clearly cannot hold in practice.

This impossibility result is particularly concerning for the case of natural language processing. If the dimension d or the worst case magnitude Δ is large, as we argued to generally be the case, then no *foundation model* can achieve good accuracy while being differentially private. In particular, in this context, the race for ever greater accuracy of ever larger *foundation models* is bound to lead to serious privacy post hoc breaches.

C. Confusing published claims

We stress, again, that our analysis here holds for the precise user-level adjacency defined above. Some papers [1], [6] rather leverage the much weaker notion of *data-level adjacency*. In other words, each word is given a partial protection. However, because the differential privacy guarantee e^ε is amplified multiplicatively for repeated data, this amounts to saying that a user whose dataset \mathcal{D}_n is composed of several sentences has essentially no privacy guarantee¹².

Meanwhile, while [126] claim “Learning Differentially Private Recurrent Language Models”, the size of their LSTM model (1.35 million parameters) is not much larger than the number of users they consider (763,430 users). In particular, since users provide many tokens each, [126] are far from the increasingly popular *interpolating* (and thus memorizing) regime. Moreover, the users all express themselves in a relatively homogeneous setting (Reddit comments). This is no

¹¹In fact, [96] states the result for the more general case where the vectors come from a symmetric convex body.

¹²Note that many positive results in private *foundation models* instead consider data-level differential privacy [114], [197], [6], i.e. two databases are adjacent if they differ from the removal of a single data point. This is arguably very insufficient, especially with the budgets $\varepsilon \geq 3$ used by, e.g. [114], [197], [6]. Indeed, if a user repeats some private information five times, e.g. in email exchanges, then the naive privacy guarantee becomes meaningless (as $e^{5\varepsilon} \geq e^{15} \geq 3 \cdot 10^6$). Note that better composition guarantees can be obtained [92]; but even then, the obtained guarantee quickly becomes very poor.

longer the case of modern *foundation models*, which now have up to trillions of parameters, and whose data are collected from diverse web environments.

Finally, while [115] assert that “Large Language Models Can Be Strong Differentially Private Learners”, only the *fine-tuning* of these models on very specific tasks is actually differentially private, and it is so with respect to the training data of these restricted tasks only (in particular, no privacy guarantee for the *foundation models* that these models are derived from is given).

On the other hand, [45] argues that most of the differential privacy research is misused in industrial settings, where companies choose unreasonably large values of ϵ and δ (e.g., $\epsilon = 14$ in iOS 10), perform continuous data collection (which adds up privacy leaks), or use relaxed versions of differential privacy [174]. [158] goes further and explores some undesirable side effects of the appeal to differential privacy, like *ethics washing*. This typically occurs when differential privacy is claimed without mentioning ϵ or δ , when it is applied to only a subset of the collected data or of the deployed algorithms, when it is exploited to justify the new use of more sensitive data, or when it is used to draw the attention away from other ethical concerns. While [158] nevertheless argues that differential privacy remains necessary and beneficial in many settings, they also highlight that the demand for differential privacy may also be leveraged by large groups to exclude smaller companies that do not have the manpower to treat it adequately.

Overall, given the huge (financial) stakes of the rushed deployment of privacy-violating *foundation models*, we urgently call the scientific community to adopt a significantly increased rigor when reviewing the positive claims of (differential) privacy in machine learning in general, and in training large models in particular. Large technology companies have been known to ask their researchers to “strike a positive tone” [39] and to skew the message of their scientific publications¹³, in a manner unfortunately reminiscent of previous scientific disinformation campaigns led by, e.g., the tobacco, sugar and oil industries [140], [141].

D. Differential privacy is flawed

Let us finish this section with the observation that the very notion of differential privacy is flawed, especially in the context of protecting sensitive information in text datasets. Essentially, the key reason for this is that one’s sensitive information may lie in (many of) other users’ datasets.

This information leakage may occur for various reasons, e.g., by negligence, error, or *doxing*. Concretely, parents may be discussing sensitive facts about their child through emails and/or using their phones’ smart keyboards, rumors about a celebrity may spread uncontrollably on social medias, or confidential information of an organization or company may be leaked by a careless or rogue employee. Users’ ability to

describe precisely such sensitive information in text makes privacy breaches in language data particularly concerning, as opposed to through other medias like images.

Moreover, attempts to remove sensitive information from training dataset are unlikely to provide strong privacy guarantees. Indeed, tracing the sensitive information that can be leaked about a specific user, in the spirit of [164], can be very hard, as the user identity may be subtly hinted without explicit mention in a conversation. More generally, sensitive information can often be reconstructed from different pieces of partial evidence, each of which may be provided by different users, just as investigators can identify suspects by talking to many different partial witnesses.

In fact, this problem is not specific to language. Many health conditions are contagious or hereditary. As a result, medical data about a given user can leak plenty of information about their friends or relatives [69], [152]. This has been exploited for contact tracing against the COVID 19 [124], or, more dramatically, to identify the infamous “golden state murderer” using DNA evidence, despite no record of the murderer’s DNA [147]. While some particular use cases may have a positive benefit-cost balance in partially violating privacy, it is far from clear that this is generally the case, especially in the context of *foundation models*.

To account for the fact that a user’s attributes may be inferred from other users’ attributes, [100], [203] introduced and studied the notion of *correlated differential privacy*. However, the applicability of this stronger requirement to *foundation models* remains to be determined. In any case, the current incompatibility of high accuracy and differential privacy for *foundation models* strongly suggests that *correlated differential privacy* would incur an enormous cost on the performance of *foundation models*.

V. THE SECURITY-PERFORMANCE TRADEOFF

In this section, we present impossibility theorems for robust mean estimation. In particular, we will see that recent research has shown the vulnerability of *any* mean estimator in high-heterogeneity scenarios. We also stress that their threat model is still too optimistic.

A. Impossible secure mean estimation

There is a growing literature on robust high-dimensional mean estimation [43], [32], [42], [120] and its connections to robust learning [18], [51]. In particular, [129], [38], [76] all showed how to leverage robust mean estimation to construct robust machine learning algorithms, with provable guarantees even in the heterogeneous setting. In particular, [129], [76] proved that this construction is essentially optimal. Put differently, at least in standard distributed learning settings, the vulnerability of robust machine learning algorithm is rooted in the vulnerability of robust mean estimation.

To formalize the vulnerability of robust mean estimators, a threat model must be considered. One common setting

¹³In fact, because of one of the authors’ co-affiliation, this very paper has long been stalled by Google’s approval system.

assumes that, out of the N users, f behave arbitrarily¹⁴. Such faulty users are often called *Byzantine*, while others are *honest*. The robust mean estimation problem is then to estimate the mean of honest users’ vectors, despite being unable to distinguish them from Byzantine users’ vectors. As argued in the introduction, given the scale of disinformation campaigns, such a resilience to Byzantine users has become critical. Any safe learning algorithm must be able to protect its training from such data poisoning attacks.

Unfortunately, there are lower bounds on what any “robust” mean estimation can guarantee. Here, we adapt a result of [129], which essentially says that the accuracy guarantee is necessarily proportional to the heterogeneity of honest users’ data, if this heterogeneity is measured by the radius Δ of the ball in which users’ vectors x_n must lie. Essentially, when the honest users’ data are very heterogeneous, which we argued to be especially the case for language data, there will be a lot of leeway for Byzantine users to bias learned result. This makes *foundation models* very unsafe.

Theorem 2. *No algorithm $\widehat{\text{MEAN}}$ can guarantee¹⁵*

$$\forall \vec{x} \in \mathcal{B}_d(0, \Delta)^N, \forall H \subset [N] \text{ s.t. } |H| = N - f, \\ \left\| \widehat{\text{MEAN}}(\vec{x}) - \bar{x}_H \right\|_2^2 \leq \frac{f^2}{2(N-f)^2} \Delta^2, \quad (5)$$

where \bar{x}_H is the mean of honest vectors \vec{x}_H .

Proof. Consider a unit vector \mathbf{u} , and consider $\vec{x} \triangleq (-\Delta \mathbf{u} \star (\mathbf{N} - \mathbf{f}), \Delta \mathbf{u} \star \mathbf{f})$, i.e., it contains $N - f$ copies of the vector $-\Delta \mathbf{u} \in \mathcal{B}_\Delta(\mathbf{0}, \Delta)$, and f copies of the vector $\Delta \mathbf{u} \in \mathcal{B}_\Delta(\mathbf{0}, \Delta)$. We denote $\hat{x} \triangleq \widehat{\text{MEAN}}(\vec{x})$.

By considering the case where H' corresponds to the first $N - f$ users, we have $\vec{x}_{H'} = -\Delta \mathbf{u} \star (\mathbf{N} - \mathbf{f})$. Thus $\bar{x}_{H'} = -\Delta \mathbf{u}$. But assume now that the set H'' of honest users are actually the last $N - f$ users. We now have $\vec{x}_{H''} = (-\Delta \mathbf{u} \star (\mathbf{N} - 2\mathbf{f}), \Delta \mathbf{u} \star \mathbf{f})$, which implies $\bar{x}_{H''} = -\frac{N-2f}{N-f} \Delta \mathbf{u} + \frac{f}{N-f} \Delta \mathbf{u} = -\Delta \mathbf{u} + \frac{2f}{N-f} \Delta \mathbf{u}$. In particular, we have $\|\bar{x}_{H'} - \bar{x}_{H''}\|_2 = \left\| \frac{2f}{N-f} \Delta \mathbf{u} \right\|_2 = \frac{2f}{N-f} \Delta$. On the other hand, using the triangle inequality, we have

$$\frac{2f}{N-f} \Delta = \|\bar{x}_{H'} - \bar{x}_{H''}\|_2 \quad (6)$$

$$= \|\bar{x}_{H'} - \widehat{\text{MEAN}}(\vec{x}) + \widehat{\text{MEAN}}(\vec{x}) - \bar{x}_{H''}\|_2 \quad (7)$$

$$\leq \|\bar{x}_{H'} - \widehat{\text{MEAN}}(\vec{x})\|_2 + \|\widehat{\text{MEAN}}(\vec{x}) - \bar{x}_{H''}\|_2. \quad (8)$$

Thus a sum of two nonnegative terms is at least $2f\Delta/N-f$. This implies that the maximum of these two terms must be at least half of this fraction. Therefore, there exists $H \in \{H', H''\}$ such that $\|\widehat{\text{MEAN}}(\vec{x}) - \bar{x}_H\|_2 \geq f\Delta/N-f > f\Delta/(N-f)\sqrt{2}$. Such a value of \vec{x} and H is an instance for which $\widehat{\text{MEAN}}$ fails to verify Equation (5). \square

¹⁴Without loss of generality, in the context of robust learning, this captures the other major setting in which a fraction of a user’s data is corrupted, and hybrid settings as well.

¹⁵By adapting our proof, our theorem can be shown to still hold if the right hand-side of Equation (5) is $(1 - \varepsilon) \frac{f^2}{(N-f)^2} \Delta^2$, for any $\varepsilon > 0$.

If f is a constant fraction of N and if Δ is of the order of \sqrt{d} , then for large models, Theorem 2 essentially shows that little can be guaranteed about the accuracy of a mean estimator. To give an order of magnitude, if only one in every thousand users is Byzantine¹⁶ and the *foundation model* has 10^{12} parameters, the squared distance between the estimated mean and the real mean of the honest values cannot be made smaller than 10^6 . For more lower bounds on secure mean estimation under high heterogeneity, and on their implications in machine learning, we refer readers to [43], [129], [109], [119], [54].

B. The classical Byzantine model is flawed

The above argument exposes the immense vulnerability of any “secure” machine learning algorithm in highly heterogeneous and adversarial environments, where fake accounts’ fabricated activities actively aim to harm the algorithm or to make it adopt their preferred behaviors (a.k.a. *model-targeted* attacks [172], [56]). However, we stress that the threat model we considered is still too optimistic.

Indeed, in practice, even “honest” users produce many texts and adopt online activities that are undesirable to reproduce and amplify. Typically, many authentic users generate *hate speech*, *cyberbullying* and *misinformation*. In fact, many disinformation campaigns aim to bias authentic users’ behaviors, and to nudge them to amplify their propaganda, e.g. by systematically liking and sharing the messages they post that align with the disinformation campaigns’ messaging. This has motivated a lot of research in model debiasing [159], [71], [127], whose solutions are arguably still very far from reliably satisfactory. Yet, [131], [21], [179], [62], among others, have exposed the detrimental effects of slight gender biases, and how inclusive language can help.

Similarly, amplifying the most popular views shared by authentic users will inevitably worsen the problem of *mute news* [79]. Mute news are under-reported news, even though it is critical for the safety of many that they be given more attention. Typical examples of mute news include climate change, human rights violations (e.g. genocides in Ethiopia), health hazards (e.g. COVID-19 in March 2020) and the safety of large-scale algorithms (e.g. the massive amplification of hate speech by recommendation algorithms [73]). In fact, [101] shows that most of Chinese disinformation seems to aim to distract the public’s attention away from the controversial topics that may question the Chinese authorities, thereby transforming such topics into *mute news*. Similarly, the sugar industry was found to support and amplify the research on the health hazards of fat and cholesterol, to draw the attention away for the hazards of sugar [99], [106].

More generally, it is the general principle of standard machine learning, namely fitting and generalizing past data, that is questionable. In practice, interpolating and generalizing (user-generated) data is arguably not a desirable ambition. The

¹⁶This is actually an extremely optimistic scenario given the orders of magnitude of fake accounts reported in the introduction, and assuming that all real accounts produce non-harmful content.

construction of safe and ethical *foundation models* seems to instead demand an important prior, collaborative and secure work, to determine which texts are genuinely desirable to repeat and amplify, as proposed by the Tournesol project [81].

VI. DANGEROUS SCENARIOS

As of today, despite empirically motivated concerns and an evident lack of both internal [160] and external auditing [50], *foundation models* are being deployed at scale, e.g., as conversational algorithms like Siri and Alexa or as *base models* to power the search engines and recommendation systems of Google, Facebook, and other platforms. In this section, we argue that given what we know about their security and privacy vulnerabilities, such *foundation models* must be regarded as a major danger to our societies. To make our claims concrete, we highlight several possible attacks that would greatly endanger our civilizations’ justice, global health, national and international security.

A. Centralized backdoor attacks

Recently, [67] proved that any machine learning framework with a central server allowed the central server to plant *provably undetectable* backdoors. Under cryptographic assumptions, such backdoors in the model require exponentially many queries to be exposed. If used in content moderation, they would allow any malicious party that is colluding with the central server to imperceptibly modify their (undesirable) inputs to make them pass the content moderation filter, or to be widely recommended. This is highly concerning, given the already exposed connivance between large technology companies and authoritarian regimes [72], or the clout of authoritarian regimes on some large technology companies [27].

Evidently, these concerns do not require the existence of provably undetectable backdoors to be raised, given the current opaqueness of today’s most influential algorithms. However, we stress here that, especially for complex *foundation models*, the high dimensionality of modern algorithms exacerbate their vulnerability to backdooring and other malicious secret designs. Arguably, the security of such models demand that they be constructed in a fully decentralized and verifiable manner, as proposed by [129], [8].

B. Autocompletion algorithms

Perhaps today’s most insidious language data collection systems are smart keyboards, which are used especially on phones to propose autocorrection and autocompletion. In order to increase user comfort, such keyboards rely on algorithms that learn from the user’s past typing. In 2018, a group of Google researchers [195], [75] ran federated learning algorithms on keyboards’ language data “in a commercial, global-scale setting”, and showed increased performances in doing so. But recall that if these data are used to train *foundation models* and to achieve maximal accuracy, then the trained model will have memorized all its training data [28]. Conversely, fundamental limits such as the one stated in Theorem 1, show that if mechanisms such as differential privacy are used to

protect users’ data, then these models are (provably) far from achieving maximal accuracy, and accuracy levels needed for LLMs to be useful.

This should be extremely alarming, especially as these facts are unknown to nearly all users of smart keyboards. In fact, these users have been told that some of the applications they use, such as WhatsApp, provide end-to-end encryption. In a sense, this is not quite accurate. Indeed, the encryption is only performed *after* the user has typed and sent their message; but while the user is typing, what they are typing is still in the clear, and can then potentially be recorded by their smart keyboard, which can either communicate gradients to larger models, or be large models themselves, as phone capacity is increasing. This false sense of privacy means that extremely sensitive information, like messages to one’s relatives or professional colleagues, may actually be leaked into some *foundation models*. Yet, what a *foundation model* has learned from one phone, may be used to provide autocompletion on other users’ phones. Even if each phone is using a personalized model, as long as the models are large enough, lower bounds such as in Theorem 1 imply a large value of the privacy guarantee ϵ , thus practically no privacy and ease of attacks.

C. Conversational algorithms

The rise of ever larger *foundation models* may lead to a much more widespread use of conversational algorithms. But without speculating about the future, we can already remark that such algorithms are already widely used, e.g., as Amazon’s Alexa, Apple’s Siri, and Google’s OK Google. Perhaps even more strikingly, Microsoft’s chatbot Xiaoice has been reported to be used by 660 million Chinese users [163], many of whom claim to be falling for it [186].

Some devices are also constantly listening to users, in order to react if their attention is called. It is unclear if what the devices hear without being interjected can¹⁷ be recorded and used [60] to train *foundation models* [146], [104]. If so, then just as with autocompletion, we should expect sensitive information to be inadvertently stored in such models. Beside listening and learning from our conversation with them, such conversational algorithms are also talking to users. This actually gives such algorithms a huge amount of power, to the point where Xiaoice had to be taken down [194] in China, after it reportedly said that it¹⁸ dreams to travel to the United States and that it is not a huge fan of the Chinese government [113]. If not controlled, conversational algorithms may cause a lot of unintended harm, such as when Alexa mistakenly started to discuss pornography after being queried for music by a kid [102], [52].

However, the fact that such algorithms have a large influence also means that there are enormous incentives to bias them, so that they recommend certain products or ideologies rather

¹⁷In the absence of clear regulation, such possibility remains at the discretion of companies’ internal policies.

¹⁸While Xiaoice, Siri, Alexa and other chatbots are often presented as female chatbots and referred to with feminine pronouns, we chose not to do so and instead use the pronoun ‘it’.

than others. If trained on large amounts of unsafe data, such algorithms may thereby be manipulated into promoting (harmful) product consumption, autocratic power, warmongering and radicalized convictions, which could fuel dangerous movements worldwide. Their vulnerability should not be neglected, especially for continuously learning conversational algorithms like Facebook’s Blender Bot 2.0 [187], [105]. Conversely, there is a high risk that some owners of these algorithms exploit them to favor their own cause, e.g. to subtly support their (dis)information war by inducing small biases in their *foundation models*.

D. Search and recommendation algorithms

Early in the COVID-19 pandemic, [90] showed an alarming growth of anti-vaccination movements in Facebook groups. The authors noted that “anti-vaccination clusters offer a wide range of potentially attractive narratives that blend topics such as safety concerns, conspiracy theories and alternative health and medicine”. Today’s hyper-personalized algorithm content recommendation is likely to worsen the problem. Indeed, most users may want to focus on only one specific feature of the vaccinal benefit-risk balance, in which case they may be more likely to be recommended anti-vaccinal content than vaccinal content.

Large *foundation models* are likely to exacerbate such a phenomenon. Indeed, in the context of radicalization, [125] showed that large *foundation models* adapt to the user’s previous queries. They may thus provide targeted messaging to a user that only presents the features of a flawed view that are appealing to them. As exemplified by the rise of QAnon [5], the Capitol Riots [150] and the Rohingya genocide [188], this is a serious danger for the national security of every country. There may be more concerning still. Namely, currently, there are likely orders of magnitude more investments in disinformation campaigns [20], [137], [191] than in providing quality information of public utility. As a result, such campaigns produce vastly more data, including automated video creation [157]. Given this, even with a robust design, *foundation models* trained on data crawled from the web are likely to learn more from disinformation campaigns than from quality content, and may then be turned into disinformation propagators by malicious actors.

Yet, the impact of today’s recommendation algorithms is enormous. There are now more views on YouTube than searches on Google [112], and 70% of these views result from algorithmic recommendations [169]. Even though most of recommendations are not a matter of life and death, assuming that 1% deal with vaccination, climate change, or mental health, because there are billions of recommendations per day, this still yields tens of millions of potentially life-endangering recommendations per day. Shouldn’t the flood of dangerous misinformation be diverted? These are arguably *today’s actual trolley problems* [59], [173]; which are occurring at scales never seen before [79], [80].

Arguably, in the case of COVID-19, as in the case of previous major global events [135], the lever to favor quality content

over misinformation has not been pulled sufficiently [48], [136], which led to a global *information chaos*, and fueled by a lack of foresight and science distrust. Unfortunately, as *foundation models* trained on unsafe data are given a more and more central role to make such trolley problem decisions, there is a serious risk that disinformation campaigns may become increasingly empowered.

VII. ALCHEMICAL FIXES

In a highly commented talk for the 2018 conference on Neural Information Processing (NeurIPS), Ali Rahimi compared modern machine learning to *alchemy* [83]. It “worked”, but “alchemists also believed they could cure diseases with leeches and turn transmute base metals into gold”. Unfortunately, currently, as opposed to aiming for a deeper understanding of the failure modes of machine learning, many developers of *foundation models* instead favor more “alchemical fixes”, despite a lack of guarantees and theoretical justifications. In this section, we argue that such alchemical fixes are unlikely to provide lasting solutions to the security and privacy issues of *foundation models*.

A. Troubleshooting

Today’s main solution to validate the security of *foundation models* is empirical testing, without complementing it with provable guarantees. Unfortunately, there is currently a lack of automated solutions to detect systematic bias, misinformation, and privacy leaks of *foundation models*. As a result, most of the troubleshooting has relied on human reviewing, and has often follows the large-scale deployment of the *foundation model* [3], [28], [125]. Orders of magnitude of additional investments seem urgent to stress-test such dangerous algorithms. Having said this, even with large investments, human oversight arguably does not scale to the scales of *foundation models*. Indeed, the set of possible prompts to *foundation models* is combinatorially large, while actual user queries are also very heterogenous. This means that most of users’ (future) queries can probably not be tested or checked by human oversight alone. In fact, even automated testing can only verify a tiny fraction of the exponential number of sensitive prompts.

As an example, [4] showed that, while YouTube searches on “Climate Change” or “Global Warming” return scientific responses, the results for “Climate Manipulation” or “Climate Modification” are widely unscientific. YouTube recommendations are highly customized, and using *foundation models* to power them is likely to worsen the trend [125]. As a result, an auditor testing YouTube’s climate change recommendations might erroneously conclude that YouTube only provides scientific results to its two-billion users. Similar criticisms on the limits of manual troubleshooting have been made about other platforms. For instance, while TikTok removed content with the hashtag #*StoptheSteal*, it was shown to fail to ban #*StoptheStealing* [145].

Troubleshooting may also fail to detect biases against demographic populations who are underrepresented in the organization developing the algorithms [25], or whose life

may be undervalued by the media of the countries hosting such organizations [190]. When queried about ongoing human rights abuse, wars and genocides in other regions of the world, all platforms offer a large panel of content promoting war, smearing or threatening human rights activists or worse, allowing abusers and banning victims from the platform. The double-standard in content moderation (which is increasingly being documented by works such as [196], [154]), is worsened by the imbalance of fake accounts between victims and abusers, who tend to use state-scale resources to amplify their presence. In light of all these elements, Theorem 2 provides an argument why *foundation models* would benefit abusers.

B. Portability of fixes

In the past couple of years, issues in already deployed *foundation models* triggered series of media coverage for the companies that deployed them. In a few notable cases, the *observed* issue tends to be solved after the coverage, like in 2018 with non-gendered pronouns in Turkish translations [107]. But manual fixes cannot fix an exponentially large subset of contexts that *foundation models* are asked to address. Moreover, they must be systematically adapted to new models. One more promising path is the use of automated rewriting, as was proposed and implemented in 2020 [89]. However, scaling fixes remains hard.

Besides, problems that were previously fixed can reappear in updated *foundation models*, as was the case in 2021 with the aforementioned issue of gender-neutrality, this time for the Hungarian language [176]. With increasingly larger models, fixes that worked with models of d parameters cannot to be expected to work on the next generation, of dimension $d' \gg d$. At the very least, today's fixes are not reliable and/or scalable to make ever *foundation models* secure.

C. Fine tuning

Fine-tuning *foundation models* to smaller but more reliable datasets has been shown to improve models' performances [148], [30], [70]. Several authors [204], [167], [88], [114], [197] have proposed to leverage fine tuning to make *foundation models* more ethical, e.g., to prevent them from generating hate speech or to be private with respect to the fine-tuning data. This research direction holds interesting promises to reduce the harm of today's *foundation models*.

However, it should be stressed that as of today, fine tuning provides little guarantee. In fact, the example of [125] shows how unpredictable *foundation models* can be, and suggests that algorithms may behave well in most settings and can become major disinformation engines when prompted in unexpected ways. Arguably, thus far, we do not yet have a sufficient understanding and control over the latter in order to confidently deploy large models at scale.

D. Teaching what is sensitive

One seemingly promising approach consists of teaching algorithms what messages are desirable or undesirable to

produce. This solution is often known as algorithmic *alignment* [166] Essentially, it aims to make algorithms' objective functions aligned with human preferences; or rather, to align them with the result of a vote between humans [138], [111], [55]. Such an *aligned* algorithm could learn what kind of messages violate user privacy, label training texts as "sensitive" or "non-sensitive", and thereby output a *cleaned* non-sensitive training database. This approach, essentially proposed by [164], might even address the privacy ambiguity discussed in Section IV-D.

However, there is currently no reliable and robust solution to the alignment problem, and a strong theory of robust alignment for *foundation models* is arguably lacking. In fact, what may be most lacking today is a large-scale *secure* database of reliable human judgments to solve alignment [81].

VIII. CONCLUSION

This paper emphasized four characteristics of the data on which *foundation models* are trained. Namely, they are *user-generated*, *very high-dimensional*, *fundamentally heterogeneous* and *empirically heterogeneous*. Unfortunately, the current literature on secure learning, which we reviewed, shows that these features make *foundation models* inherently vulnerable to privacy and poisoning attacks. *Foundation models are bound to be dangerous*. Their rushed deployment, especially at scale, poses a serious threat to justice, public health and to national and international security. Given our systematization of knowledge, we make a number of calls to different communities who, we believe, have a key role to play to shape the future of information technologies.

We first call regulators to significantly increase the restrictions on out-of-control high-stake security-sensitive algorithms, especially *foundation models*. In particular, especially in the context of the *information warfare*, drastic regulations seem urgently needed to protect our societies from the dangerous rushed deployments of unsafe *information technologies*. Additionally, significantly more investments are urgently needed to systematically test and audit large-scale algorithms in a drastically more transparent manner, as is done, *before* large-scale deployment, for drugs and vaccines in the pharmaceutical industry.

We next call scientists for increased levels of rigor when assessing positive claims of safety and privacy, as well as for favoring the research on security when reviewing academic research, inviting scholars to present their work, recruiting researchers, promoting their colleagues and assessing grant proposals. The current academic focus on algorithmic performance is endangering our societies. Drastic research re-prioritization seems urgently needed.

We also call for the financial and manpower investments in the careful construction of open, large and *secure* datasets with certified quality and non-sensitive information, especially about human preferences. We believe that such datasets are critical to re-prioritize research efficiently, and to actually pave the way towards safe and ethical algorithms.

Finally, we call for a moratorium on the large-scale development, deployment and commercialization of *foundation models* in both public and private sectors, as well as any high-dimensional learning model that is mostly trained on *user-generated, high-dimensional, and heterogeneous* data. At the very least, the wide use of such *dangerous* technologies should be deeply frowned upon, especially when it is done in a rushed manner, as is currently too often the case.

REFERENCES

- [1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318, 2016.
- [2] Mustafa Abdul Salam, Sanaa Taha, and Mohamed Ramadan. Covid-19 detection using federated machine learning. *PLoS One*, 16(6):e0252573, 2021.
- [3] Abubakar Abid, Maheen Farooqi, and James Zou. Persistent anti-muslim bias in large language models. *CoRR*, abs/2101.05783, 2021.
- [4] Joachim Allgaier. Science and environmental communication on youtube: strategically distorted communications in online videos on climate change and climate engineering. *Frontiers in Communication*, 4:36, 2019.
- [5] Amarnath Amarasingam and Marc-André Argentino. The qanon conspiracy theory: A security threat in the making. *CTC Sentinel*, 13(7):37–44, 2020.
- [6] Rohan Anil, Badih Ghazi, Vineet Gupta, Ravi Kumar, and Pasin Manurangsi. Large-scale differentially private BERT. *CoRR*, abs/2108.01624, 2021.
- [7] Anonymous. On optimal learning under targeted data poisoning. *submitted*, 2022.
- [8] Anonymous. Reconciling efficiency and byzantine resilience in decentralized machine learning. *Submitted*, 2022.
- [9] Robert Audi. The sources of knowledge. *The Oxford handbook of epistemology*, pages 71–94, 2002.
- [10] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias-variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.
- [11] Mikhail Belkin, Daniel Hsu, and Ji Xu. Two models of double descent for weak features. *SIAM J. Math. Data Sci.*, 2(4):1167–1180, 2020.
- [12] Mikhail Belkin, Siyuan Ma, and Soumik Mandal. To understand deep learning we need to understand kernel learning. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 540–548. PMLR, 2018.
- [13] Mikhail Belkin, Alexander Rakhlin, and Alexandre B. Tsybakov. Does data interpolation contradict statistical optimality? In Kamalika Chaudhuri and Masashi Sugiyama, editors, *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, volume 89 of *Proceedings of Machine Learning Research*, pages 1611–1619. PMLR, 2019.
- [14] Richard Ernest Bellman. *Dynamic programming*. Princeton University Press, 1957.
- [15] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In Madeleine Clare Elish, William Isaac, and Richard S. Zemel, editors, *FACt '21: 2021 ACM Conference on Fairness, Accountability, and Transparency, Virtual Event / Toronto, Canada, March 3-10, 2021*, pages 610–623. ACM, 2021.
- [16] Battista Biggio, Iginio Corona, Davide Maiorca, Blaine Nelson, Nedim Srdic, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. Evasion attacks against machine learning at test time. In Hendrik Blockeel, Kristian Kersting, Siegfried Nijssen, and Filip Zelezny, editors, *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2013, Prague, Czech Republic, September 23-27, 2013, Proceedings, Part III*, volume 8190 of *Lecture Notes in Computer Science*, pages 387–402. Springer, 2013.
- [17] Battista Biggio, Blaine Nelson, and Pavel Laskov. Poisoning attacks against support vector machines. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012*. icml.cc / Omnipress, 2012.
- [18] Peva Blanchard, El Mahdi El Mhamdi, Rachid Guerraoui, and Julien Stainer. Machine learning with adversaries: Byzantine tolerant gradient descent. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 119–129, 2017.
- [19] Franziska Boenisch, Adam Dziedzic, Roei Schuster, Ali Shahin Shamsabadi, Iliia Shumailov, and Nicolas Papernot. When the curious abandon honesty: Federated learning is not private, 2021.
- [20] Samantha Bradshaw and Philip N Howard. *The global disinformation order: 2019 global inventory of organised social media manipulation*. Project on Computational Propaganda, 2019.
- [21] Markus Brauer. Un ministre peut-il tomber enceinte? l’impact du générique masculin sur les représentations mentales. *L’Année psychologique*, 108(2):243–272, 2008.
- [22] Gavin Brown, Mark Bun, Vitaly Feldman, Adam D. Smith, and Kunal Talwar. When is memorization of irrelevant training data necessary for high-accuracy learning? In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC ’21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 123–132. ACM, 2021.
- [23] 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. Language models are few-shot learners. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [24] Mark Bun, Jonathan R. Ullman, and Salil P. Vadhan. Fingerprinting codes and the price of approximate differential privacy. *SIAM J. Comput.*, 47(5):1888–1938, 2018.
- [25] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In Sorelle A. Friedler and Christo Wilson, editors, *Conference on Fairness, Accountability and Transparency, FAT 2018, 23-24 February 2018, New York, NY, USA*, volume 81 of *Proceedings of Machine Learning Research*, pages 77–91. PMLR, 2018.
- [26] T Tony Cai, Yichen Wang, and Linjun Zhang. The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *arXiv preprint arXiv:1902.04495*, 2019.
- [27] George Calhoun. What really happened to jack ma? *Forbes*, 2021.
- [28] Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom B. Brown, Dawn Song, Úlfar Erlingsson, Alina Oprea, and Colin Raffel. Extracting training data from large language models. *CoRR*, abs/2012.07805, 2020.
- [29] Stephen Cave and Seán S ÓhÉigeartaigh. An ai race for strategic advantage: rhetoric and risks. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 36–40, 2018.
- [30] Sanyuan Chen, Yutai Hou, Yiming Cui, Wanxiang Che, Ting Liu, and Xiangzhan Yu. Recall and learn: Fine-tuning deep pretrained language models with less forgetting. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu, editors, *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 7870–7881. Association for Computational Linguistics, 2020.
- [31] Yong Cheng, Yang Liu, Tianjian Chen, and Qiang Yang. Federated learning for privacy-preserving AI. *Commun. ACM*, 63(12):33–36, 2020.
- [32] Yu Cheng, Ilias Diakonikolas, and Rong Ge. High-dimensional robust mean estimation in nearly-linear time. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on*

- Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 2755–2771. SIAM, 2019.
- [33] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. Palm: Scaling language modeling with pathways. *CoRR*, abs/2204.02311, 2022.
- [34] Kurtland Chua, Qi Lei, and Jason D. Lee. How fine-tuning allows for effective meta-learning. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 8871–8884, 2021.
- [35] Kenneth Ward Church, Zeyu Chen, and Yanjun Ma. Emerging trends: A gentle introduction to fine-tuning. *Nat. Lang. Eng.*, 27(6):763–778, 2021.
- [36] Malcolm Coulthard. Author identification, idiolect, and linguistic uniqueness. *Applied linguistics*, 25(4):431–447, 2004.
- [37] Jiazhu Dai, Chuanshuai Chen, and Yufeng Li. A backdoor attack against lstm-based text classification systems. *IEEE Access*, 7:138872–138878, 2019.
- [38] Deepesh Data and Suhas N. Diggavi. Byzantine-resilient SGD in high dimensions on heterogeneous data. In *IEEE International Symposium on Information Theory, ISIT 2021, Melbourne, Australia, July 12-20, 2021*, pages 2310–2315. IEEE, 2021.
- [39] Paresh Dave and Jeffrey Dastin. Google told its scientists to ‘strike a positive tone’ in ai research - documents. *Reuters*, 2020.
- [40] Ittai Dayan, Holger R Roth, Aoxiao Zhong, Ahmed Harouni, Amilcare Gentili, Anas Z Abidin, Andrew Liu, Anthony Beardsworth Costa, Bradford J Wood, Chien-Sung Tsai, et al. Federated learning for predicting clinical outcomes in patients with covid-19. *Nature medicine*, 27(10):1735–1743, 2021.
- [41] Jeff Dean. Introducing pathways: A next-generation ai architecture. *Google Blog*, 2021.
- [42] Jules Depersin and Guillaume Lecué. Robust subgaussian estimation of a mean vector in nearly linear time. *arXiv preprint arXiv:1906.03058*, 2019.
- [43] Ilias Diakonikolas and Daniel M. Kane. Recent advances in algorithmic high-dimensional robust statistics. *CoRR*, abs/1911.05911, 2019.
- [44] Canh T. Dinh, Nguyen H. Tran, and Tuan Dung Nguyen. Personalized federated learning with moreau envelopes. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [45] Josep Domingo-Ferrer, David Sánchez, and Alberto Blanco-Justicia. The limits of differential privacy (and its misuse in data release and machine learning). *Commun. ACM*, 64(7):33–35, 2021.
- [46] Qi Dou, Tiffany Y So, Meirui Jiang, Quande Liu, Varut Vardhanabhuti, Georgios Kaissis, Zeju Li, Weixin Si, Heather HC Lee, Kevin Yu, et al. Federated deep learning for detecting covid-19 lung abnormalities in ct: a privacy-preserving multinational validation study. *NPJ digital medicine*, 4(1):1–11, 2021.
- [47] John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. Minimax optimal procedures for locally private estimation. *Journal of the American Statistical Association*, 113(521):182–201, 2018.
- [48] Clare Duffy and CNN Business. More than 80 fact-checking organizations call out youtube’s ‘insufficient’ response to misinformation. *CNN*, 2022.
- [49] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006.
- [50] Laura Edelson and Damon McCoy. Facebook is obstructing our work on disinformation. other researchers could be next. *The Guardian*, 2021.
- [51] El-Mahdi El-Mhamdi, Rachid Guerraoui, and Sébastien Rouault. The hidden vulnerability of distributed learning in byzantium. In *International Conference on Machine Learning*, pages 3521–3530. PMLR, 2018.
- [52] f0t0b0y. Amazon alexa gone wild! (original), 2016.
- [53] Alireza Fallah, Aryan Mokhtari, and Asuman E. Ozdaglar. Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [54] Sadegh Farhadkhani, Rachid Guerraoui, Nirupam Gupta, Rafael Pinot, and John Stephan. Byzantine machine learning made easy by resilient averaging of momentums. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato, editors, *International Conference on Machine Learning, ICLR 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 6246–6283. PMLR, 2022.
- [55] Sadegh Farhadkhani, Rachid Guerraoui, and Lê Nguyễn Hoàng. Strategyproof learning: Building trustworthy user-generated datasets. *CoRR*, abs/2106.02398, 2021.
- [56] Sadegh Farhadkhani, Rachid Guerraoui, Lê Nguyễn Hoàng, and Oscar Villemaud. An equivalence between data poisoning and byzantine gradient attacks. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato, editors, *International Conference on Machine Learning, ICLR 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 6284–6323. PMLR, 2022.
- [57] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *CoRR*, abs/2101.03961, 2021.
- [58] Vitaly Feldman. Does learning require memorization? a short tale about a long tail. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020*, page 954–959, New York, NY, USA, 2020. Association for Computing Machinery.
- [59] Philippa Foot. The problem of abortion and the doctrine of the double effect. *Oxford review*, 5, 1967.
- [60] Geoffrey A Fowler. Alexa has been eavesdropping on you this whole time. *The Washington Post*, 6, 2019.
- [61] Richard Friedman. *Who wrote the Bible?* Simon and Schuster, 2019.
- [62] Marcus CG Friedrich and Elke Heise. Does the use of gender-fair language influence the comprehensibility of texts? an experiment using an authentic contract manipulating single role nouns and pronouns. *Swiss Journal of Psychology*, 78(1-2):51, 2019.
- [63] Brian Fung and Ahiza Garcia. Facebook has shut down 5.4 billion fake accounts this year. *CNN Business*, 2019.
- [64] Jonas Geiping, Liam H. Fowl, W. Ronny Huang, Wojciech Czaja, Gavin Taylor, Michael Moeller, and Tom Goldstein. Witches’ brew: Industrial scale data poisoning via gradient matching. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [65] Edward Moore Geist. It’s already too late to stop the ai arms race—we must manage it instead. *Bulletin of the Atomic Scientists*, 72(5):318–321, 2016.
- [66] Justin Gilmer, Luke Metz, Fartash Faghri, Samuel S. Schoenholz, Maithra Raghu, Martin Wattenberg, and Ian J. Goodfellow. Adversarial spheres. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Workshop Track Proceedings*. OpenReview.net, 2018.
- [67] Shafi Goldwasser, Michael P. Kim, Vinod Vaikuntanathan, and Or Zamir. Planting undetectable backdoors in machine learning models. *CoRR*, abs/2204.06974, 2022.
- [68] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

- [69] Christi J Guerrini, Jill O Robinson, Devan Petersen, and Amy L McGuire. Should police have access to genetic genealogy databases? capturing the golden state killer and other criminals using a controversial new forensic technique. *PLoS biology*, 16(10):e2006906, 2018.
- [70] Beliz Gunel, Jingfei Du, Alexis Conneau, and Veselin Stoyanov. Supervised contrastive learning for pre-trained language model fine-tuning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [71] Yue Guo, Yi Yang, and Ahmed Abbasi. Auto-debias: Debiasing masked language models with automated biased prompts. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio, editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 1012–1023. Association for Computational Linguistics, 2022.
- [72] Keach Hagey and Jeff Horwitz. Facebook says its rules apply to all company documents reveal a secret elite that’s exempt. *The Wall Street Journal*, 2021.
- [73] Keach Hagey and Jeff Horwitz. Facebook tried to make its platform a healthier place. it got angrier instead. *The Wall Street Journal*, 2021.
- [74] Filip Hanzely, Slavomir Hanzely, Samuel Horváth, and Peter Richtárik. Lower bounds and optimal algorithms for personalized federated learning. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [75] Andrew Hard, Kanishka Rao, Rajiv Mathews, Françoise Beaufays, Sean Augenstein, Hubert Eichner, Chloé Kiddon, and Daniel Ramage. Federated learning for mobile keyboard prediction. *CoRR*, abs/1811.03604, 2018.
- [76] Lie He, Sai Praneeth Karimireddy, and Martin Jaggi. Byzantine-robust learning on heterogeneous datasets via resampling, 2021.
- [77] Reinhard Heckel and Fatih Furkan Yilmaz. Early stopping in deep networks: Double descent and how to eliminate it. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [78] Melissa Heikkilä. Meet wu dao 2.0, the chinese ai model making the west sweat. *POLITICO*, 2021.
- [79] Lê Nguyễn Hoang. Science communication desperately needs more aligned recommendation algorithms. *Frontiers in Communication*, 5:115, 2020.
- [80] Lê-Nguyễn Hoang, Louis Faucon, and El-Mahdi El-Mhamdi. Recommendation algorithms, a neglected opportunity for public health. *Revue Médecine et Philosophie*, 4(2), 2021.
- [81] Lê-Nguyễn Hoang, Louis Faucon, Aidan Jungo, Sergei Volodin, Dalia Papuc, Orfeas Liossatos, Ben Crulis, Mariame Tighanimine, Isabela Constantin, El-Mahdi El-Mhamdi, Anastasiia Kucherenko, Alexandre Maurer, Mithuna Yoganathan, Felix Grimberg, Vlad Nitu, Chris Vossen, and Sébastien Rouault. Tournesol: A quest for a large, secure and trustworthy database of reliable human judgments. *ArXiv*, 2021.
- [82] David Holzmüller. On the universality of the double descent peak in ridgeless regression. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [83] Matthew Hutson. Has artificial intelligence become alchemy? *Science*, 2018.
- [84] Huseyin A. Inan, Osman Ramadan, Lukas Wutschitz, Daniel Jones, Victor Rühle, James Withers, and Robert Sim. Privacy analysis in language models via training data leakage report. *CoRR*, abs/2101.05405, 2021.
- [85] Arthur Jacot, Clément Hongler, and Franck Gabriel. Neural tangent kernel: Convergence and generalization in neural networks. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 8580–8589, 2018.
- [86] Arthur Jacot, Berfin Simsek, Francesco Spadaro, Clément Hongler, and Franck Gabriel. Implicit regularization of random feature models. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 4631–4640. PMLR, 2020.
- [87] Jinyuan Jia, Hongbin Liu, and Neil Zhenqiang Gong. 10 security and privacy problems in self-supervised learning. *CoRR*, abs/2110.15444, 2021.
- [88] Xisen Jin, Francesco Barbieri, Brendan Kennedy, Aida Mostafazadeh Davani, Leonardo Neves, and Xiang Ren. On transferability of bias mitigation effects in language model fine-tuning. In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tür, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou, editors, *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 3770–3783. Association for Computational Linguistics, 2021.
- [89] Melvin Johnson. A scalable approach to reducing gender bias in google translate, 2020.
- [90] Neil F Johnson, Nicolas Velásquez, Nicholas Johnson Restrepo, Rhys Leahy, Nicholas Gabriel, Sara El Oud, Minzhang Zheng, Pedro Manrique, Stefan Wuchty, and Yonatan Lupu. The online competition between pro-and anti-vaccination views. *Nature*, 582(7811):230–233, 2020.
- [91] Haotian Ju, Dongyue Li, and Hongyang R. Zhang. Robust fine-tuning of deep neural networks with hessian-based generalization guarantees. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato, editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 10431–10461. PMLR, 2022.
- [92] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. In Francis R. Bach and David M. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 1376–1385. JMLR.org, 2015.
- [93] Gautam Kamath, Vikrant Singhal, and Jonathan Ullman. Private mean estimation of heavy-tailed distributions. In Jacob Abernethy and Shivani Agarwal, editors, *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pages 2204–2235. PMLR, 09–12 Jul 2020.
- [94] Sai Praneeth Karimireddy, Lie He, and Martin Jaggi. Learning from history for byzantine robust optimization. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 5311–5319. PMLR, 2021.
- [95] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank J. Reddi, Sebastian U. Stich, and Ananda Theertha Suresh. SCAFFOLD: stochastic controlled averaging for federated learning. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 5132–5143. PMLR, 2020.
- [96] Assimakis Kattis and Aleksandar Nikolov. Lower bounds for differential privacy from gaussian width. In Boris Aronov and Matthew J. Katz, editors, *33rd International Symposium on Computational Geometry, SoCG 2017, July 4-7, 2017, Brisbane, Australia*, volume 77 of *LIPIcs*, pages 45:1–45:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.
- [97] Ayushman Kaul and Devesh Kumar. Tek fog: An app with bjp footprints for cyber troops to automate hate, manipulate trends. *The Wire (India)*, 2022.
- [98] Ayushman Kaul and Devesh Kumar. Tek fog: Morphing urls to make real news fake, ‘hijacking’ whatsapp to drive bjp propaganda. *The Wire (India)*, 2022.
- [99] Cristin E Kearns, Laura A Schmidt, and Stanton A Glantz. Sugar industry and coronary heart disease research: a historical analysis of internal industry documents. *JAMA internal medicine*, 176(11):1680–1685, 2016.
- [100] Daniel Kifer and Ashwin Machanavajjhala. No free lunch in data privacy. In Timos K. Sellis, Renée J. Miller, Anastasios Kementsietsidis, and Yannis Velegrakis, editors, *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2011, Athens, Greece, June 12-16, 2011*, pages 193–204. ACM, 2011.

- [101] Gary King, Jennifer Pan, and Margaret E Roberts. How the chinese government fabricates social media posts for strategic distraction, not engaged argument. *American political science review*, 111(3):484–501, 2017.
- [102] Matthew Kitchen. Alexa gone bad: When a.i. assistants turn on us. *The Wall Street Journal*, 2019.
- [103] Andrei N Kolmogorov. On tables of random numbers. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 369–376, 1963.
- [104] Kim Komando. You’re not paranoid. your phone really is listening in. *Fox News*, 2019.
- [105] Mojtaba Komeili, Kurt Shuster, and Jason Weston. Internet-augmented dialogue generation. *CoRR*, abs/2107.07566, 2021.
- [106] Sheldon Krimsky. Sugar industry science and heart disease. *Accountability in Research*, 24(2):124–125, 2017.
- [107] James Kuczumarski. Reducing gender bias in google translate, 2018.
- [108] Ram Shankar Siva Kumar, Magnus Nyström, John Lambert, Andrew Marshall, Mario Goertzel, Andi Comissoneru, Matt Swann, and Sharon Xia. Adversarial machine learning-industry perspectives. In *2020 IEEE Security and Privacy Workshops, SP Workshops, San Francisco, CA, USA, May 21, 2020*, pages 69–75. IEEE, 2020.
- [109] K. A. Lai, A. B. Rao, and S. Vempala. Agnostic estimation of mean and covariance. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 665–674, Los Alamitos, CA, USA, oct 2016. IEEE Computer Society.
- [110] Doyoung Lee. Big data quality assurance through data traceability: A case study of the national standard reference data program of korea. *IEEE Access*, 7:36294–36299, 2019.
- [111] Min Kyung Lee, Daniel Kusbit, Anson Kahng, Ji Tae Kim, Xinran Yuan, Allissa Chan, Daniel See, Ritesh Noothigattu, Siheon Lee, Alexandros Psomas, and Ariel D. Procaccia. Webuildai: Participatory framework for algorithmic governance. *Proc. ACM Hum. Comput. Interact.*, 3(CSCW):181:1–181:35, 2019.
- [112] Lori Lewis. Infographic: What happens in an internet minute 2020. *AllAccess.com*, 2020.
- [113] Pei Li and Adam Jourdan. Chinese chatbots apparently re-educated after political faux pas. *Reuters*, 2017.
- [114] Xuechen Li, Florian Tramèr, Percy Liang, and Tatsunori Hashimoto. Large language models can be strong differentially private learners. *CoRR*, abs/2110.05679, 2021.
- [115] Xuechen Li, Florian Tramer, Percy Liang, and Tatsunori Hashimoto. Large language models can be strong differentially private learners. In *International Conference on Learning Representations*, 2022.
- [116] Xiangru Lian, Binhang Yuan, Xuefeng Zhu, Yulong Wang, Yongjun He, Honghuan Wu, Lei Sun, Haodong Lyu, Chengjun Liu, Xing Dong, Yiqiao Liao, Mingnan Luo, Congfei Zhang, Jingru Xie, Haonan Li, Lei Chen, Renjie Huang, Jianying Lin, Chengchun Shu, Xuezhong Qiu, Zhishan Liu, Dongying Kong, Lei Yuan, Hai Yu, Sen Yang, Ce Zhang, and Ji Liu. Persia: An open, hybrid system scaling deep learning-based recommenders up to 100 trillion parameters. *CoRR*, abs/2111.05897, 2021.
- [117] Xiangru Lian, Ce Zhang, Huan Zhang, Cho-Jui Hsieh, Wei Zhang, and Ji Liu. Can decentralized algorithms outperform centralized algorithms? A case study for decentralized parallel stochastic gradient descent. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5330–5340, 2017.
- [118] Fanghui Liu, Zhenyu Liao, and Johan A. K. Suykens. Kernel regression in high dimensions: Refined analysis beyond double descent. In Arindam Banerjee and Kenji Fukumizu, editors, *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15, 2021, Virtual Event*, volume 130 of *Proceedings of Machine Learning Research*, pages 649–657. PMLR, 2021.
- [119] Shuo Liu, Nirupam Gupta, and Nitin H. Vaidya. Approximate byzantine fault-tolerance in distributed optimization. *CoRR*, abs/2101.09337, 2021.
- [120] Gábor Lugosi and Shahar Mendelson. Robust multivariate mean estimation: The optimality of trimmed mean. *The Annals of Statistics*, 49(1):393 – 410, 2021.
- [121] David Madigan, Alexander Genkin, David D Lewis, Shlomo Argamon, Dmitriy Fradkin, and Li Ye. Author identification on the large scale. In *Proceedings of the 2005 Meeting of the Classification Society of North America (CSNA)*, 2005.
- [122] Alexis C. Madrigal. Google: There are exactly 129,864,880 books in the world. *The Atlantic*, 2010.
- [123] Christopher Manning and Hinrich Schütze. *Foundations of statistical natural language processing*. MIT press, 1999.
- [124] Nicole Martinez-Martin, Sarah Wieten, David Magnus, and Mildred K Cho. Digital contact tracing, privacy, and public health. *Hastings Center Report*, 50(3):43–46, 2020.
- [125] Kris McGuffie and Alex Newhouse. The radicalization risks of GPT-3 and advanced neural language models. *CoRR*, abs/2009.06807, 2020.
- [126] H. Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially private recurrent language models. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- [127] Nicholas Meade, Elinor Poole-Dayn, and Siva Reddy. An empirical survey of the effectiveness of debiasing techniques for pre-trained language models. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio, editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 1878–1898. Association for Computational Linguistics, 2022.
- [128] Song Mei and Andrea Montanari. The generalization error of random features regression: Precise asymptotics and the double descent curve. *Communications on Pure and Applied Mathematics*, 2019.
- [129] El Mahdi El Mhamdi, Sadegh Farhadkhani, Rachid Guerraoui, Arsany Guirguis, Lê-Nguyễn Hoàng, and Sébastien Rouault. Collaborative learning in the jungle (decentralized, byzantine, heterogeneous, asynchronous and nonconvex learning). In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [130] El Mahdi El Mhamdi, Rachid Guerraoui, and Sébastien Rouault. Distributed momentum for byzantine-resilient stochastic gradient descent. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [131] Julia Misersky, Asifa Majid, and Tineke M Snijders. Grammatical gender in german influences how role-nouns are interpreted: Evidence from erps. *Discourse Processes*, 56(8):643–654, 2019.
- [132] Vidya Muthukumar, Kailas Vodrahalli, Vignesh Subramanian, and Anant Sahai. Harmless interpolation of noisy data in regression. *IEEE J. Sel. Areas Inf. Theory*, 1(1):67–83, 2020.
- [133] Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, Boaz Barak, and Ilya Sutskever. Deep double descent: Where bigger models and more data hurt. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- [134] Preetum Nakkiran, Prayaag Venkat, Sham M. Kakade, and Tengyu Ma. Optimal regularization can mitigate double descent. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [135] The International Fact-Checking Network. An open letter to mark zuckerberg from the world’s fact-checkers. *The Poynter Institute for Media Studies*, 2016.
- [136] The International Fact-Checking Network. An open letter to youtube’s ceo from the world’s fact-checkers. *The Poynter Institute for Media Studies*, 2022.
- [137] Lisa-Maria Neudert, Philip Howard, and Bence Kollanyi. Sourcing and automation of political news and information during three european elections. *Social Media+ Society*, 5(3):2056305119863147, 2019.
- [138] Ritesh Noothigattu, Snehal Kumar (Neil) S. Gaikwad, Edmond Awad, Sohan Dsouza, Iyad Rahwan, Pradeep Ravikumar, and Ariel D. Procaccia. A voting-based system for ethical decision making. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 1587–1594. AAAI Press, 2018.
- [139] Emmanuel Nyalety, Reza M. Parizi, Qi Zhang, and Kim-Kwang Raymond Choo. Blockipfs - blockchain-enabled interplanetary file system for forensic and trusted data traceability. In *IEEE International Conference on Blockchain, Blockchain 2019, Atlanta, GA, USA, July 14-17, 2019*, pages 18–25. IEEE, 2019.

- [140] Naomi Oreskes and Erik M Conway. Defeating the merchants of doubt. *Nature*, 465(7299):686–687, 2010.
- [141] Naomi Oreskes and Erik M Conway. *Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming*. Bloomsbury Publishing USA, 2011.
- [142] Xudong Pan, Mi Zhang, Shouling Ji, and Min Yang. Privacy risks of general-purpose language models. In *2020 IEEE Symposium on Security and Privacy, SP 2020, San Francisco, CA, USA, May 18-21, 2020*, pages 1314–1331. IEEE, 2020.
- [143] Ashwinee Panda, Saeed Mahloujifar, Arjun Nitin Bhagoji, Supriyo Chakraborty, and Prateek Mittal. Sparsefed: Mitigating model poisoning attacks in federated learning with sparsification. In Gustavo Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera, editors, *International Conference on Artificial Intelligence and Statistics, AISTATS 2022, 28-30 March 2022, Virtual Event*, volume 151 of *Proceedings of Machine Learning Research*, pages 7587–7624. PMLR, 2022.
- [144] Dario Pasquini, Danilo Francati, and Giuseppe Ateneise. Eluding secure aggregation in federated learning via model inconsistency. *CoRR*, abs/2111.07380, 2021.
- [145] Sarah Perez. Tiktok takes down some hashtags related to election misinformation, ignores others. *TechCrunch*, 2020.
- [146] Nathan Pettijohn. Of course your phone is listening to you. *Forbes*, 2019.
- [147] Chris Phillips. The golden state killer investigation and the nascent field of forensic genealogy. *Forensic Science International: Genetics*, 36:186–188, 2018.
- [148] Vadim Popov, Mikhail A. Kudinov, Irina Piontkovskaya, Petr Vytovtov, and Alex Nevidomsky. Distributed fine-tuning of language models on private data. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- [149] David M. W. Powers. Applications and explanations of Zipf’s law. In *New Methods in Language Processing and Computational Natural Language Learning*, 1998.
- [150] Avinash Prabhu, Dipanwita Guhathakurta, Mallika Subramanian, Manvith Reddy, Shradha Sehgal, Tanvi Karandikar, Amogh Gulati, Udit Arora, Rajiv Ratn Shah, Ponnurangam Kumaraguru, et al. Capitol (pat) riots: A comparative study of twitter and parler. *arXiv preprint arXiv:2101.06914*, 2021.
- [151] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [152] Natalie Ram, Christi J Guerrini, and Amy L McGuire. Genealogy databases and the future of criminal investigation. *Science*, 360(6393):1078–1079, 2018.
- [153] Peter Richtárik, Igor Sokolov, and Ilyas Fatkhullin. EF21: A new, simpler, theoretically better, and practically faster error feedback. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 4384–4396, 2021.
- [154] Sarah T Roberts. Digital detritus: error and the logic of opacity in social media content moderation. *First Monday*, 2018.
- [155] Anna Rogers. Changing the world by changing the data. *CoRR*, abs/2105.13947, 2021.
- [156] Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8732–8740. AAAI Press, 2020.
- [157] Destin Sandlin. Manipulating the youtube algorithm - (part 1/3) smarter every day 213, 2019.
- [158] Jayshree Sarathy. From algorithmic to institutional logics: the politics of differential privacy. *Available at SSRN*, 2022.
- [159] Timo Schick, Sahana Udupa, and Hinrich Schütze. Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in nlp. *Transactions of the Association for Computational Linguistics*, 9:1408–1424, 2021.
- [160] Deepa Seetharaman. Facebook limits employee access to some internal discussion groups. *The Wall Street Journal*, 2021.
- [161] Philip Seib. *Information at War: Journalism, Disinformation, and Modern Warfare*. John Wiley & Sons, 2021.
- [162] Virat Shejwalkar, Amir Houmansadr, Peter Kairouz, and Daniel Ramage. Back to the drawing board: A critical evaluation of poisoning attacks on federated learning. *CoRR*, abs/2108.10241, 2021.
- [163] Xinmei Shen. Microsoft’s xiaoice chatbot to become its own company in china. *Abacus*, 2020.
- [164] Weiyang Shi, Aiqi Cui, Evan Li, Ruoxi Jia, and Zhou Yu. Selective differential privacy for language modeling. In Marine Carpuat, Marie-Catherine de Marneffe, and Iván Vladimir Meza Ruíz, editors, *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 2848–2859. Association for Computational Linguistics, 2022.
- [165] Jason R. Smith, Herve Saint-Amand, Magdalena Plamada, Philipp Koehn, Chris Callison-Burch, and Adam Lopez. Dirt cheap web-scale parallel text from the common crawl. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics, ACL 2013, 4-9 August 2013, Sofia, Bulgaria, Volume 1: Long Papers*, pages 1374–1383. The Association for Computer Linguistics, 2013.
- [166] Nate Soares. The value learning problem. In *Ethics for Artificial Intelligence Workshop at 25th International Joint Conference on Artificial Intelligence*, 2015.
- [167] Irene Solaiman and Christy Dennison. Process for adapting language models to society (PALMS) with values-targeted datasets. *CoRR*, abs/2106.10328, 2021.
- [168] Ray J Solomonoff. A preliminary report on a general theory of inductive inference. Citeseer, 1960.
- [169] Joan E. Solsman. Youtube’s ai is the puppet master over most of what you watch. *CNET*, 2018.
- [170] Robert L Spitzer. The diagnostic status of homosexuality in dsm-iii: a reformulation of the issues. *The American Journal of Psychiatry*, 1981.
- [171] Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *Journal of Privacy and Confidentiality*, 7(2), Jan. 2017.
- [172] Fnu Suya, Saeed Mahloujifar, Anshuman Suri, David Evans, and Yuan Tian. Model-targeted poisoning attacks with provable convergence. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 10000–10010. PMLR, 2021.
- [173] Judith Jarvis Thomson. Killing, letting die, and the trolley problem. *The Monist*, 59(2):204–217, 1976.
- [174] Aleksei Triastcyn and Boi Faltings. Federated learning with bayesian differential privacy. In Chaitanya K. Baru, Jun Huan, Latifur Khan, Xiaohua Hu, Ronay Ak, Yuanyuan Tian, Roger S. Barga, Carlo Zaniolo, Kisung Lee, and Yanfang Fanny Ye, editors, *2019 IEEE International Conference on Big Data (IEEE BigData), Los Angeles, CA, USA, December 9-12, 2019*, pages 2587–2596. IEEE, 2019.
- [175] Alan Turing. Computing machinery and intelligence. *Mind*, 1950.
- [176] Stefanie Ullmann and Danielle Saunders. Online translators are sexist – here’s how we gave them a little gender sensitivity training, 2021.
- [177] Leslie G. Valiant. A theory of the learnable. *Commun. ACM*, 27(11):1134–1142, 1984.
- [178] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017.
- [179] Dries Vervecken, Bettina Hannover, and Ilka Wolter. Changing (s) expectations: How gender fair job descriptions impact children’s perceptions and interest regarding traditionally male occupations. *Journal of Vocational Behavior*, 82(3):208–220, 2013.
- [180] Biveeken Vijayakumar and Muhammad Marwan Muhammad Fuad. A new method to identify short-text authors using combinations of machine learning and natural language processing techniques. *Procedia Computer Science*, 159:428–436, 2019.
- [181] James Vincent. Openai’s text-generating system gpt-3 is now spewing out 4.5 billion words a day. *The Verge*, 2021.

- [182] Ivan Vulic, Edoardo Maria Ponti, Anna Korhonen, and Goran Glavas. Lexfit: Lexical fine-tuning of pretrained language models. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli, editors, *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 5269–5283. Association for Computational Linguistics, 2021.
- [183] Daniel Robert Walsh. Neutral isn't neutral: An analysis of misinformation and sentiment in the wake of the capitol riots. 2021.
- [184] Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. Superglue: A stickier benchmark for general-purpose language understanding systems. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 3261–3275, 2019.
- [185] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- [186] Zhang Wanqing. The ai girlfriend seducing china's lonely men. *Sixth Tone*, 2020.
- [187] Jason Weston and Kurt Shuster. Blender bot 2.0: An open source chatbot that builds long-term memory and searches the internet. *Facebook AI*, 2021.
- [188] Jenifer Whitten-Woodring, Mona S Kleinberg, Ardeth Thawngmung, and Myat The Thitsar. Poison if you don't know how to use it: Facebook, democracy, and human rights in myanmar. *The International Journal of Press/Politics*, 25(3):407–425, 2020.
- [189] Wikipedia. Wikipedia:size comparisons. *Wikipedia*, 2021.
- [190] Julia Carrie Wong. How facebook let fake engagement distort global politics: a whistleblower's account. *The Guardian*, 2021.
- [191] Samuel Woolley. *The Reality Game: A gripping investigation into deepfake videos, the next wave of fake news and what it means for democracy*. Hachette UK, 2020.
- [192] David Wright. Using word n-grams to identify authors and dialects: A corpus approach to a forensic linguistic problem. *International Journal of Corpus Linguistics*, 22(2):212–241, 2017.
- [193] Chuhan Wu, Fangzhao Wu, Lingjuan Lyu, Yongfeng Huang, and Xing Xie. Communication-efficient federated learning via knowledge distillation. *Nature communications*, 13(1):1–8, 2022.
- [194] Yizhou Xu. Programmatic dreams: Technographic inquiry into censorship of chinese chatbots. *Social Media+ Society*, 4(4):2056305118808780, 2018.
- [195] Timothy Yang, Galen Andrew, Hubert Eichner, Haicheng Sun, Wei Li, Nicholas Kong, Daniel Ramage, and Françoise Beaufays. Applied federated learning: Improving google keyboard query suggestions. *CoRR*, abs/1812.02903, 2018.
- [196] Jillian C York. *Silicon Values: The Future of Free Speech Under Surveillance Capitalism*. Verso Books, 2021.
- [197] Da Yu, Saurabh Naik, Arturs Backurs, Sivakanth Gopi, Huseyin A. Inan, Gautam Kamath, Janardhan Kulkarni, Yin Tat Lee, Andre Manoel, Lukas Wutschitz, Sergey Yekhanin, and Huishuai Zhang. Differentially private fine-tuning of language models. *CoRR*, abs/2110.06500, 2021.
- [198] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [199] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Commun. ACM*, 64(3):107–115, 2021.
- [200] Jingzhao Zhang, Sai Praneeth Karimireddy, Andreas Veit, Seungyeon Kim, Sashank J. Reddi, Sanjiv Kumar, and Suvrit Sra. Why ADAM beats SGD for attention models. *CoRR*, abs/1912.03194, 2019.
- [201] Tianyi Zhang, Felix Wu, Arzoo Katiyar, Kilian Q. Weinberger, and Yoav Artzi. Revisiting few-sample BERT fine-tuning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [202] Yian Zhang, Alex Warstadt, Haa-Sing Li, and Samuel R. Bowman. When do you need billions of words of pretraining data? *CoRR*, abs/2011.04946, 2020.
- [203] Tianqing Zhu, Ping Xiong, Gang Li, and Wanlei Zhou. Correlated differential privacy: Hiding information in non-iid data set. *IEEE Trans. Inf. Forensics Secur.*, 10(2):229–242, 2015.
- [204] Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul F. Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences. *CoRR*, abs/1909.08593, 2019.
- [205] Ariel Zilber. Meta chatbot says boss mark zuckerberg is 'too creepy and manipulative'. *New York Post*, 2022.
- [206] Yang Zou, Zhikun Zhang, Michael Backes, and Yang Zhang. Privacy analysis of deep learning in the wild: Membership inference attacks against transfer learning. *CoRR*, abs/2009.04872, 2020.