

Advances and Open Problems in Federated Learning

Other titles in Foundations and Trends® in Machine Learning

Graph Kernels: State-of-the-Art and Future Challenges

Karsten Borgwardt, Elisabetta Ghisu, Felipe Llinares-López, Leslie O'Bray and Bastian Rieck

ISBN: 978-1-68083-770-4

Data Analytics on Graphs Part III: Machine Learning on Graphs, from Graph Topology to Applications

Ljubiša Stanković, Danilo Mandić, Miloš Daković, Miloš Brajović, Bruno Scalzo, Shengxi Li and Anthony G. Constantinides

ISBN: 978-1-68083-982-16

Data Analytics on Graphs Part II: Signals on Graphs

Ljubiša Stanković, Danilo Mandić, Miloš Daković, Miloš Brajović, Bruno Scalzo, Shengxi Li and Anthony G. Constantinides

ISBN: 978-1-68083-982-1

Data Analytics on Graphs Part I: Graphs and Spectra on Graphs

Ljubiša Stanković, Danilo Mandić, Miloš Daković, Miloš Brajović, Bruno Scalzo, Shengxi Li and Anthony G. Constantinides

ISBN: 978-1-68083-982-1

Advances and Open Problems in Federated Learning

Peter Kairouz
Google Research
Kairouz@google.com

H. Brendan McMahan
Google Research

et al.

now

the essence of knowledge

Boston — Delft

Foundations and Trends[®] in Machine Learning

Published, sold and distributed by:

now Publishers Inc.
PO Box 1024
Hanover, MA 02339
United States
Tel. +1-781-985-4510
www.nowpublishers.com
sales@nowpublishers.com

Outside North America:

now Publishers Inc.
PO Box 179
2600 AD Delft
The Netherlands
Tel. +31-6-51115274

The preferred citation for this publication is

Peter Kairouz, H. Brendan McMahan *et al.*. *Advances and Open Problems in Federated Learning*. Foundations and Trends[®] in Machine Learning, vol. 14, no. 1–2, pp. 1–210, 2021.

ISBN: 978-1-68083-789-6

© 2021 Peter Kairouz, H. Brendan McMahan *et al.*

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: www.copyright.com

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1 781 871 0245; www.nowpublishers.com; sales@nowpublishers.com

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, www.nowpublishers.com; e-mail: sales@nowpublishers.com

Foundations and Trends[®] in Machine Learning

Volume 14, Issue 1–2, 2021

Editorial Board

Editor-in-Chief

Michael Jordan

University of California, Berkeley
United States

Editors

Peter Bartlett
UC Berkeley

Yoshua Bengio
Université de Montréal

Avrim Blum
*Toyota Technological
Institute*

Craig Boutilier
University of Toronto

Stephen Boyd
Stanford University

Carla Brodley
Northeastern University

Inderjit Dhillon
Texas at Austin

Jerome Friedman
Stanford University

Kenji Fukumizu
ISM

Zoubin Ghahramani
Cambridge University

David Heckerman
Amazon

Tom Heskes
Radboud University

Geoffrey Hinton
University of Toronto

Aapo Hyvarinen
Helsinki IIT

Leslie Pack Kaelbling
MIT

Michael Kearns
UPenn

Daphne Koller
Stanford University

John Lafferty
Yale

Michael Littman
Brown University

Gabor Lugosi
Pompeu Fabra

David Madigan
Columbia University

Pascal Massart
Université de Paris-Sud

Andrew McCallum
*University of
Massachusetts Amherst*

Marina Meila
University of Washington

Andrew Moore
CMU

John Platt
Microsoft Research

Luc de Raedt
KU Leuven

Christian Robert
Paris-Dauphine

Sunita Sarawagi
IIT Bombay

Robert Schapire
Microsoft Research

Bernhard Schoelkopf
Max Planck Institute

Richard Sutton
University of Alberta

Larry Wasserman
CMU

Bin Yu
UC Berkeley

Editorial Scope

Topics

Foundations and Trends® in Machine Learning publishes survey and tutorial articles in the following topics:

- Adaptive control and signal processing
- Applications and case studies
- Behavioral, cognitive and neural learning
- Bayesian learning
- Classification and prediction
- Clustering
- Data mining
- Dimensionality reduction
- Evaluation
- Game theoretic learning
- Graphical models
- Independent component analysis
- Inductive logic programming
- Kernel methods
- Markov chain Monte Carlo
- Model choice
- Nonparametric methods
- Online learning
- Optimization
- Reinforcement learning
- Relational learning
- Robustness
- Spectral methods
- Statistical learning theory
- Variational inference
- Visualization

Information for Librarians

Foundations and Trends® in Machine Learning, 2021, Volume 14, 6 issues. ISSN paper version 1935-8237. ISSN online version 1935-8245. Also available as a combined paper and online subscription.

Contents

1	Introduction	4
1.1	The Cross-Device Federated Learning Setting	7
1.2	Federated Learning Research	14
1.3	Organization	15
2	Relaxing the Core FL Assumptions: Applications to Emerging Settings and Scenarios	16
2.1	Fully Decentralized/Peer-to-Peer Distributed Learning	16
2.2	Cross-Silo Federated Learning	23
2.3	Split Learning	26
2.4	Executive Summary	28
3	Improving Efficiency and Effectiveness	29
3.1	Non-IID Data in Federated Learning	29
3.2	Optimization Algorithms for Federated Learning	33
3.3	Multi-Task Learning, Personalization, and Meta-Learning	46
3.4	Adapting ML Workflows for Federated Learning	51
3.5	Communication and Compression	53
3.6	Application to More Types of Machine Learning Problems and Models	57
3.7	Executive Summary	58

4	Preserving the Privacy of User Data	60
4.1	Actors, Threat Models, and Privacy in Depth	63
4.2	Tools and Technologies	65
4.3	Protections Against External Malicious Actors	81
4.4	Protections Against an Adversarial Server	89
4.5	User Perception	99
4.6	Executive Summary	102
5	Defending Against Attacks and Failures	104
5.1	Adversarial Attacks on Model Performance	105
5.2	Non-Malicious Failure Modes	120
5.3	Exploring the Tension Between Privacy and Robustness	123
5.4	Executive Summary	124
6	Ensuring Fairness and Addressing Sources of Bias	126
6.1	Bias in Training Data	127
6.2	Fairness Without Access to Sensitive Attributes	128
6.3	Fairness, Privacy, and Robustness	130
6.4	Leveraging Federation to Improve Model Diversity	132
6.5	Federated Fairness: New Opportunities and Challenges	133
6.6	Executive Summary	135
7	Addressing System Challenges	136
7.1	Platform Development and Deployment Challenges	136
7.2	System Induced Bias	138
7.3	System Parameter Tuning	143
7.4	On-Device Runtime	145
7.5	The Cross-Silo Setting	147
7.6	Executive Summary	148
8	Concluding Remarks	150
	Acknowledgments	152

Appendices	153
A.1 Software and Datasets for Federated Learning	154
References	159

Advances and Open Problems in Federated Learning

Peter Kairouz^{*1}, H. Brendan McMahan^{*2} *et al.*

¹*Google Research, USA; Kairouz@google.com*

²*Google Research, USA*

ABSTRACT

Federated learning (FL) is a machine learning setting where many clients (e.g., mobile devices or whole organizations) collaboratively train a model under the orchestration of a central server (e.g., service provider), while keeping the training data decentralized. FL embodies the principles of focused data collection and minimization, and can mitigate many of the systemic privacy risks and costs resulting from traditional, centralized machine learning and data science approaches. Motivated by the explosive growth in FL research, this monograph discusses recent advances and presents an extensive collection of open problems and challenges.

Peter Kairouz and H. Brendan McMahan conceived, coordinated, and edited this work.

Peter Kairouz, H. Brendan McMahan, *et al.* (2021), “Advances and Open Problems in Federated Learning”, Foundations and Trends[®] in Machine Learning: Vol. 14, No. 1–2, pp 1–210. DOI: 10.1561/22000000083.

Full Author List

Peter Kairouz
Google Research

H. Brendan McMahan
Google Research

Brendan Avent
USC

Aurélien Bellet
INRIA

Mehdi Bennis
University of Oulu

Arjun Nitin Bhagoji
Princeton University

Kallista Bonawitz
Google Research

Zachary Charles
Google Research

Graham Cormode
University of Warwick

Rachel Cummings
Georgia Tech.

Rafael G. L. D'Oliveira
Rutgers University

Hubert Eichner
Google Research

Salim El Rouayheb
Rutgers University

David Evans
University of Virginia

Josh Gardner
University of Washington

Zachary Garrett
Google Research

Adrià Gascón
Google Research

Badih Ghazi
Google Research

Phillip B. Gibbons
CMU

Marco Gruteser
Google Research
Rutgers University

Zaid Harchaoui
University of Washington

Chaoyang He
USC

Lie He
EPFL

Zhouyuan Huo
University of Pittsburgh

Ben Hutchinson
Google Research

Justin Hsu
UW-Madison

Martin Jaggi
EPFL

Tara Javidi
UC San Diego

Gauri Joshi
CMU

Mikhail Khodak
CMU

Jakub Konečný
Google Research

Aleksandra Korolova
USC

Farinaz Koushanfar
UC San Diego

Sanmi Koyejo
Google Research
UIUC

Tancrede Lepoint
Google Research

Yang Liu
NTU

Prateek Mittal
Princeton

Mehryar Mohri
Google Research

Richard Nock
ANU

Ayfer Özgür
Stanford

Rasmus Pagh
Google Research
IT University of Copenhagen

Hang Qi
Google Research

Daniel Ramage
Google Research

Ramesh Raskar
MIT

Mariana Raykova
Google Research

Dawn Song
UC Berkeley

Weikang Song
Google Research

Sebastian U. Stich
EPFL

Ziteng Sun
Cornell

Ananda Theertha Suresh
Google Research

Florian Tramèr
Stanford

Praneeth Vepakomma
MIT

Jianyu Wang
CMU

Li Xiong
Emory

Zheng Xu
Google Research

Qiang Yang
HKUST

Felix X. Yu
Google Research

Han Yu
NTU

Sen Zhao
Google Research

1

Introduction

Federated learning (FL) is a machine learning setting where many clients (e.g., mobile devices or whole organizations) collaboratively train a model under the orchestration of a central server (e.g., service provider), while keeping the training data decentralized. It embodies the principles of focused collection and data minimization, and can mitigate many of the systemic privacy risks and costs resulting from traditional, centralized machine learning. This area has received significant interest recently, both from research and applied perspectives. This monograph describes the defining characteristics and challenges of the federated learning setting, highlights important practical constraints and considerations, and then enumerates a range of valuable research directions. The goals of this work are to highlight research problems that are of significant theoretical and practical interest, and to encourage research on problems that could have significant real-world impact.

The term *federated learning* was introduced in 2016 by McMahan *et al.* [1]: “We term our approach Federated Learning, since the learning task is solved by a loose federation of participating devices (which we refer to as clients) which are coordinated by a central server.” An unbalanced and non-IID (identically and independently distributed)

data partitioning across a massive number of unreliable devices with limited communication bandwidth was introduced as the defining set of challenges.

Significant related work predates the introduction of the term federated learning. A longstanding goal pursued by many research communities (including cryptography, databases, and machine learning) is to analyze and learn from data distributed among many owners without exposing that data. Cryptographic methods for computing on encrypted data were developed starting in the early 1980s [2], [3], and Agrawal and Srikant [4] and Vaidya *et al.* [5] are early examples of work that sought to learn from local data using a centralized server while preserving privacy. Conversely, even since the introduction of the term federated learning, we are aware of no single work that directly addresses the full set of FL challenges. Thus, the term federated learning provides a convenient shorthand for a set of characteristics, constraints, and challenges that often co-occur in applied ML problems on decentralized data where privacy is paramount.

This monograph originated at the Workshop on Federated Learning and Analytics held June 17–18th, 2019, hosted at Google’s Seattle office. During the course of this two-day event, the need for a broad paper surveying the many open challenges in the area of federated learning became clear.¹

A key property of many of the problems discussed is that they are inherently interdisciplinary—solving them likely requires not just machine learning, but techniques from distributed optimization, cryptography, security, differential privacy, fairness, compressed sensing, systems, information theory, statistics, and more. Many of the hardest problems are at the intersections of these areas, and so we believe collaboration will be essential to ongoing progress. One of the goals of this work is to highlight the ways in which techniques from these fields can potentially be combined, raising both interesting possibilities as well as new challenges.

¹During the preparation of this work, Li *et al.* [6] independently released an excellent but less comprehensive survey.

Since the term federated learning was initially introduced with an emphasis on mobile and edge device applications [1], [7], interest in applying FL to other applications has greatly increased, including some which might involve only a small number of relatively reliable clients, for example multiple organizations collaborating to train a model. We term these two federated learning settings “cross-device” and “cross-silo” respectively. Given these variations, we propose a somewhat broader definition of federated learning:

Federated learning is a machine learning setting where multiple entities (clients) collaborate in solving a machine learning problem, under the coordination of a central server or service provider. Each client’s raw data is stored locally and not exchanged or transferred; instead, focused updates intended for immediate aggregation are used to achieve the learning objective.

Focused updates are updates narrowly scoped to contain the minimum information necessary for the specific learning task at hand; aggregation is performed as early as possible in the service of data minimization. We note that this definition distinguishes federated learning from fully decentralized (peer-to-peer) learning techniques as discussed in Subsection 2.1.

Although privacy-preserving data analysis has been studied for more than 50 years, only in the past decade have solutions been widely deployed at scale (e.g., [8], [9]). Cross-device federated learning and federated data analysis are now being applied in consumer digital products. Google makes extensive use of federated learning in the Gboard mobile keyboard [10]–[14], as well as in features on Pixel phones and in Android Messages [15]. While Google has pioneered cross-device FL, interest in this setting is now much broader, for example: Apple is using cross-device FL in iOS 13 [16], for applications like the QuickType keyboard and the vocal classifier for “Hey Siri” [17]; doc.ai is developing cross-device FL solutions for medical research [18], and Snips has explored cross-device FL for hotword detection [19].

Cross-silo applications have also been proposed or described in myriad domains including finance risk prediction for reinsurance [20],

pharmaceuticals discovery [21], electronic health records mining [22], medical data segmentation [23], [24], and smart manufacturing [25].

The growing demand for federated learning technology has resulted in a number of tools and frameworks becoming available. These include TensorFlow Federated [26], Federated AI Technology Enabler [27], PySyft [28], Leaf [29], PaddleFL [30] and Clara Training Framework [31]; more details in Appendix A.1. Commercial data platforms incorporating federated learning are in development from established technology companies as well as smaller start-ups.

Table 1.1 contrasts both cross-device and cross-silo federated learning with traditional single-datacenter distributed learning across a range of axes. These characteristics establish many of the constraints that practical federated learning systems must typically satisfy, and hence serve to both motivate and inform the open challenges in federated learning. They will be discussed at length in the sections that follow.

These two FL variants are called out as representative and important examples, but different FL settings may have different combinations of these characteristics. For the remainder of this monograph, we consider the cross-device FL setting unless otherwise noted, though many of the problems apply to other FL settings as well. Section 2 specifically addresses some of the many other variations and applications.

Next, we consider cross-device federated learning in more detail, focusing on practical aspects common to a typical large-scale deployment of the technology; Bonawitz *et al.* [32] provides even more detail for a particular production system, including a discussion of specific architectural choices and considerations.

1.1 The Cross-Device Federated Learning Setting

This section takes an applied perspective, and unlike the previous section, does not attempt to be definitional. Rather, the goal is to describe some of the practical issues in cross-device FL and how they might fit into a broader machine learning development and deployment ecosystem. The hope is to provide useful context and motivation for the open problems that follow, as well as to aid researchers in estimating how straightforward it would be to deploy a particular new approach

Table 1.1: Typical characteristics of federated learning settings vs. distributed learning in the datacenter (e.g., [33]). Cross-device and cross-silo federated learning are two examples of FL domains, but are not intended to be exhaustive. The primary defining characteristics of FL are highlighted in bold, but the other characteristics are also critical in determining which techniques are applicable

	Datacenter Distributed Learning	Cross-Silo Federated Learning	Cross-Device Federated Learning
Setting	Training a model on a large but “flat” dataset. Clients are compute nodes in a single cluster or datacenter.	Training a model on siloed data. Clients are different organizations (e.g., medical or financial) or geo-distributed datacenters.	The clients are a very large number of mobile or IoT devices.
Data distribution	Data is centrally stored and can be shuffled and balanced across clients. Any client can read any part of the dataset.	Data is generated locally and remains decentralized. Each client stores its own data and cannot read the data of other clients. Data is not independently or identically distributed.	
Orchestration	Centrally orchestrated.	A central orchestration server/service organizes the training, but never sees raw data.	
Wide-area communication	None (fully connected clients in one datacenter/cluster).	Typically a hub-and-spoke topology, with the hub representing a coordinating service provider (typically without data) and the spokes connecting to clients.	
Data availability	_____ All clients are almost always available.	_____	Only a fraction of clients are available at any one time, often with diurnal or other variations.
Distribution scale	Typically 1–1000 clients.	Typically 2–100 clients.	Massively parallel, up to 10^{10} clients.
Primary bottleneck	Computation is more often the bottleneck in the datacenter, where very fast networks can be assumed.	Might be computation or communication.	Communication is often the primary bottleneck, though it depends on the task. Generally, cross-device federated computations use wi-fi or slower connections.

Continued.

Table 1.1: Continued

	Datacenter Distributed Learning	Cross-Silo Federated Learning	Cross-Device Federated Learning
Addressability	Each client has an identity or name that allows the system to access it specifically.		Clients cannot be indexed directly (i.e., no use of client identifiers).
Client statefulness	Stateful—each client may participate in each round of the computation, carrying state from round to round.		Stateless—each client will likely participate only once in a task, so generally a fresh sample of never-before-seen clients in each round of computation is assumed.
Client reliability	_____ Relatively few failures.		Highly unreliable—5% or more of the clients participating in a round of computation are expected to fail or drop out (e.g., because the device becomes ineligible when battery, network, or idleness requirements are violated).
Data partition axis	Data can be partitioned/re-partitioned arbitrarily across clients.	Partition is fixed. Could be example-partitioned (horizontal) or feature-partitioned (vertical).	Fixed partitioning by example (horizontal).




Figure 1.1: The lifecycle of an FL-trained model and the various actors in a federated learning system. This figure is revisited in Section 4 from a threat models perspective.

in a real-world system. We begin by sketching the lifecycle of a model before considering a FL training process.

1.1.1 The Lifecycle of a Model in Federated Learning

The FL process is typically driven by a model engineer developing a model for a particular application. For example, a domain expert in natural language processing may develop a next word prediction model for use in a virtual keyboard. Figure 1.1 shows the primary components and actors. At a high level, a typical workflow is:

1. **Problem identification:** The model engineer identifies a problem to be solved with FL.
2. **Client instrumentation:** If needed, the clients (e.g., an app running on mobile phones) are instrumented to store locally (with limits on time and quantity) the necessary training data. In many cases, the app already will have stored this data (e.g., a text messaging app must store text messages, a photo management app already stores photos). However, in some cases additional data or metadata might need to be maintained, e.g., user interaction data to provide labels for a supervised learning task.

3. **Simulation prototyping (optional):** The model engineer may prototype model architectures and test learning hyperparameters in an FL simulation using a proxy dataset.
4. **Federated model training:** Multiple federated training tasks are started to train different variations of the model, or use different optimization hyperparameters.
5. **(Federated) model evaluation:** After the tasks have trained sufficiently (typically a few days, see below), the models are analyzed and good candidates selected. Analysis may include metrics computed on standard datasets in the datacenter, or federated evaluation wherein the models are pushed to held-out clients for evaluation on local client data.
6. **Deployment:** Finally, once a good model is selected, it goes through a standard model launch process, including manual quality assurance, live A/B testing (usually by using the new model on some devices and the previous generation model on other devices to compare their in-vivo performance), and a staged rollout (so that poor behavior can be discovered and rolled back before affecting too many users). The specific launch process for a model is set by the owner of the application and is usually independent of how the model is trained. In other words, this step would apply equally to a model trained with federated learning or with a traditional datacenter approach.

One of the primary practical challenges an FL system faces is making the above workflow as straightforward as possible, ideally approaching the ease-of-use achieved by ML systems for centralized training. While much of this monograph concerns federated training specifically, there are many other components including federated analytics tasks like model evaluation and debugging. Improving these is the focus of Subsection 3.4. For now, we consider in more detail the training of a single FL model (Step 4 above).

1.1.2 A Typical Federated Training Process

We now consider a template for FL training that encompasses the Federated Averaging algorithm of McMahan *et al.* [1] and many others; again, variations are possible, but this gives a common starting point.

A server (service provider) orchestrates the training process, by repeating the following steps until training is stopped (at the discretion of the model engineer who is monitoring the training process):

1. **Client selection:** The server samples from a set of clients meeting eligibility requirements. For example, mobile phones might only check in to the server if they are plugged in, on an unmetered wi-fi connection, and idle, in order to avoid impacting the user of the device.
2. **Broadcast:** The selected clients download the current model weights and a training program (e.g., a TensorFlow graph [34]) from the server.
3. **Client computation:** Each selected device locally computes an update to the model by executing the training program, which might for example run SGD on the local data (as in Federated Averaging).
4. **Aggregation:** The server collects an aggregate of the device updates. For efficiency, stragglers might be dropped at this point once a sufficient number of devices have reported results. This stage is also the integration point for many other techniques which will be discussed later, possibly including: secure aggregation for added privacy, lossy compression of aggregates for communication efficiency, and noise addition and update clipping for differential privacy.
5. **Model update:** The server locally updates the shared model based on the aggregated update computed from the clients that participated in the current round.

Table 1.2 gives typical order-of-magnitude sizes for the quantities involved in a typical federated learning application on mobile devices.

Table 1.2: Order-of-magnitude sizes for typical cross-device federated learning applications

Total population size	10^6 – 10^{10} devices
Devices selected for one round of training	50–5000
Total devices that participate in training one model	10^5 – 10^7
Number of rounds for model convergence	500–10000
Wall-clock training time	1–10 days

The separation of the client computation, aggregation, and model update phases is not a strict requirement of federated learning, and it indeed excludes certain classes of algorithms, for example asynchronous SGD where each client’s update is immediately applied to the model, before any aggregation with updates from other clients. Such asynchronous approaches may simplify some aspects of system design, and also be beneficial from an optimization perspective (though this point can be debated). However, the approach presented above has a substantial advantage in affording a separation of concerns between different lines of research: advances in compression, differential privacy, and secure multi-party computation can be developed for standard primitives like computing sums or means over decentralized updates, and then composed with arbitrary optimization or analytics algorithms, so long as those algorithms are expressed in terms of aggregation primitives.

It is also worth emphasizing that in two respects, the FL training process should not impact the user experience. First, as outlined above, even though model parameters are typically sent to some devices during the broadcast phase of each round of federated training, these models are an ephemeral part of the training process, and not used to make “live” predictions shown to the user. This is crucial, because training ML models is challenging, and a misconfiguration of hyperparameters can produce a model that makes bad predictions. Instead, user-visible use of the model is deferred to a rollout process as detailed above in Step 6 of the model lifecycle. Second, the training itself is intended to be invisible to the user—as described under client selection, training does not slow the device or drain the battery because it only executes when the device is idle and connected to power. However, the limited

availability these constraints introduce leads directly to open research challenges which will be discussed subsequently, such as semi-cyclic data availability and the potential for bias in client selection.

1.2 Federated Learning Research

The remainder of this monograph surveys many open problems that are motivated by the constraints and challenges of real-world federated learning settings, from training models on medical data from a hospital system to training using hundreds of millions of mobile devices. Needless to say, most researchers working on federated learning problems will likely not be deploying production FL systems, nor have access to fleets of millions of real-world devices. This leads to a key distinction between the practical settings that motivate the work and experiments conducted in simulation which provide evidence of the suitability of a given approach to the motivating problem.

This makes FL research somewhat different than other ML fields from an experimental perspective, leading to additional considerations in conducting FL research. In particular, when highlighting open problems, we have attempted, when possible, to also indicate relevant performance metrics which can be measured in simulation, the characteristics of datasets which will make them more representative of real-world performance, etc. The need for simulation also has ramifications for the presentation of FL research. While not intended to be authoritative or absolute, we make the following modest suggestions for presenting FL research that addresses the open problems we describe:

- As shown in Table 1.1, the FL setting can encompass a wide range of problems. Compared to fields where the setting and goals are well-established, it is important to precisely describe the details of the particular FL setting of interest, particularly when the proposed approach makes assumptions that may not be appropriate in all settings (e.g., stateful clients that participate in all rounds).
- Of course, details of any simulations should be presented in order to make the research reproducible. But it is also important to

explain which aspects of the real-world setting the simulation is designed to capture (and which it is not), in order to effectively make the case that success on the simulated problem implies useful progress on the real-world objective. We hope that the guidance in this monograph will help with this.

- Privacy and communication efficiency are always first-order concerns in FL, even if the experiments are simulations running on a single machine using public data. More so than with other types of ML, for any proposed approach it is important to be unambiguous about *where computation happens* as well as *what is communicated*.

Software libraries for federated learning simulation as well as standard datasets can help ease the challenges of conducting effective FL research; Appendix A.1 summarizes some of the currently available options. Developing standard evaluation metrics and establishing standard benchmark datasets for different federated learning settings (cross-device and cross-silo) remain highly important directions for ongoing work.

1.3 Organization

Section 2 builds on the ideas in Table 1.1, exploring other FL settings and problems beyond the original focus on cross-device settings. Section 3 then turns to core questions around improving the efficiency and effectiveness of federated learning. Section 4 undertakes a careful consideration of threat models and considers a range of technologies toward the goal of achieving rigorous privacy protections. As with all machine learning systems, in federated learning applications there may be incentives to manipulate the models being trained, and failures of various kinds are inevitable; these challenges are discussed in Section 5. Finally, we address the important challenges of providing fair and unbiased models in Section 6.

References

- [1] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pp. 1273–1282, 2017.
- [2] R. L. Rivest, L. Adleman, and M. L. Dertouzos, “On data banks and privacy homomorphisms,” *Foundations of Secure Computation*, Academia Press, pp. 169–179, 1978.
- [3] A. C. Yao, “Protocols for secure computations,” in *Symposium on Foundations of Computer Science*, pp. 160–164, 1982.
- [4] R. Agrawal and R. Srikant, “Privacy-preserving data mining,” in *ACM SIGMOD International Conference on Management of Data*, pp. 439–450, 2000.
- [5] J. Vaidya, H. Yu, and X. Jiang, “Privacy-preserving SVM classification,” *Knowl. Inf. Syst.*, vol. 14, no. 2, pp. 161–178, 2008.
- [6] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, “Federated learning: Challenges, methods, and future directions,” *IEEE Signal Process. Mag.*, vol. 37, no. 3, pp. 50–60, 2020.
- [7] H. B. McMahan and D. Ramage, *Federated learning: Collaborative machine learning without centralized training data*, [Online]. Available: <https://ai.googleblog.com/2017/04/federated-learning-collaborative.html>. Google AI Blog, Apr. 2017.

- [8] Differential Privacy Team, “Learning with privacy at scale,” *Apple Machine Learning Journal*, vol. 1, no. 8, 2017. [Online]. Available: <https://machinelearning.apple.com/2017/12/06/learning-with-privacy-at-scale.html>.
- [9] Ú. Erlingsson, V. Pihur, and A. Korolova, “RAPPOR: Randomized aggregatable privacy-preserving ordinal response,” in *ACM CCS*, 2014. DOI: [10.1145/2660267.2660348](https://doi.org/10.1145/2660267.2660348).
- [10] M. Chen, R. Mathews, T. Ouyang, and F. Beaufays, “Federated learning of out-of-vocabulary words,” *arXiv preprint 1903.10635*, 2019. [Online]. Available: <http://arxiv.org/abs/1903.10635>.
- [11] A. Hard, K. Rao, R. Mathews, F. Beaufays, S. Augenstein, H. Eichner, C. Kiddon, and D. Ramage, “Federated learning for mobile keyboard prediction,” *arXiv preprint 1811.03604*, 2018.
- [12] S. Pichai, “Google’s Sundar Pichai: Privacy should not be a luxury good,” *New York Times*, May 7, 2019.
- [13] S. Ramaswamy, R. Mathews, K. Rao, and F. Beaufays, “Federated learning for emoji prediction in a mobile keyboard,” *arXiv preprint arxiv:1906.04329*, 2019.
- [14] T. Yang, G. Andrew, H. Eichner, H. Sun, W. Li, N. Kong, D. Ramage, and F. Beaufays, “Applied federated learning: Improving Google keyboard query suggestions,” *arXiv preprint arxiv:1812.02903*, 2018.
- [15] support.google, *Your chats stay private while Messages improves suggestions*, 2019. [Online]. Available: <https://support.google.com/messages/answer/9327902>. Retrieved Aug. 2019.
- [16] Apple, *Private federated learning (NeurIPS 2019 Expo Talk Abstract)*, https://nips.cc/ExpoConferences/2019/schedule?talk_id=40, 2019.
- [17] Apple, *Designing for privacy* (video and slide deck), Apple WWDC, <https://developer.apple.com/videos/play/wwdc2019/708>, 2019.
- [18] W. de Brouwer, *The federated future is ready for shipping*, https://medium.com/@_doc_ai/the-federated-future-is-ready-for-shipping-d17ff40f43e3, Mar. 2019.

- [19] D. Leroy, A. Coucke, T. Lavril, T. Gisselbrecht, and J. Dureau, “Federated learning for keyword spotting,” *arXiv preprint arXiv:1810.05512*, 2018.
- [20] WeBank, *WeBank and Swiss re signed cooperation MoU*, 2019. [Online]. Available: <https://www.fedai.org/news/webank-and-swiss-re-signed-cooperation-mou/>. Retrieved Aug. 2019.
- [21] EU CORDIS, *Machine learning ledger orchestration for drug discovery*, 2019. [Online]. Available: https://cordis.europa.eu/project/rcn/223634/factsheet/en?WT.mc_id=RSS-Feed&WT.rss_f=project&WT.rss_a=223634&WT.rss_ev=a. Retrieved Aug. 2019.
- [22] FeatureCloud, *FeatureCloud: Our vision*, 2019. [Online]. Available: <https://featurecloud.eu/about/our-vision/>. Retrieved Aug. 2019.
- [23] ai.intel, *Federated learning for medical imaging*, 2019. [Online]. Available: <https://www.intel.ai/federated-learning-for-medical-imaging/>. Retrieved Aug. 2019.
- [24] P. Courtiol, C. Maussion, M. Moarii, E. Pronier, S. Pilcer, M. Sefta, P. Manceron, S. Toldo, M. Zaslavskiy, N. Le Stang, N. Girard, O. Elemento, A. G. Nicholson, J.-Y. Blay, F. Galateau-Sallé, G. Wainrib, and T. Clozel, “Deep learning-based classification of mesothelioma improves prediction of patient outcome,” *Nature Medicine*, pp. 1–7, 2019.
- [25] Musketeer, *Musketeer: About*, 2019. [Online]. Available: <http://musketeer.eu/project/>. Retrieved Aug. 2019.
- [26] The TFF Authors, *TensorFlow Federated*, 2019. [Online]. Available: <https://www.tensorflow.org/federated>.
- [27] The FATE Authors, *Federated AI technology enabler*, 2019. [Online]. Available: <https://www.fedai.org/>.
- [28] T. Ryffel, A. Trask, M. Dahl, B. Wagner, J. Mancuso, D. Rueckert, and J. Passerat-Palmbach, “A generic framework for privacy preserving deep learning,” *arXiv preprint arXiv:1811.04017*, 2018.
- [29] The Leaf Authors, *Leaf*, 2019. [Online]. Available: <https://leaf.cmu.edu/>.

- [30] The PaddleFL Authors, *PaddleFL*, 2019. [Online]. Available: <https://github.com/PaddlePaddle/PaddleFL>.
- [31] N. Clara, *The clara training framework authors*, 2019. [Online]. Available: <https://developer.nvidia.com/clara>.
- [32] K. A. Bonawitz, H. Eichner, W. Grieskamp, D. Huba, A. Ingerman, V. Ivanov, C. M. Kiddon, J. Konečný, S. Mazzocchi, B. McMahan, T. V. Overveldt, D. Petrou, D. Ramage, and J. Roselander, “Towards federated learning at scale: System design,” in *SysML 2019*, 2019. [Online]. Available: <https://arxiv.org/abs/1902.01046>.
- [33] J. Dean, G. S. Corrado, R. Monga, K. Chen, M. Devin, Q. V. Le, M. Z. Mao, M. Ranzato, A. Senior, P. Tucker, K. Yang, and A. Y. Ng, “Large scale distributed deep networks,” in *Proceedings of the International Conference on Neural Information Processing Systems*, pp. 1223–1231, 2012.
- [34] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, *TensorFlow: Large-scale machine learning on heterogeneous systems*, Software available from tensorflow.org, 2015. [Online]. Available: <https://www.tensorflow.org/>.
- [35] P. Vanhaesebrouck, A. Bellet, and M. Tommasi, “Decentralized collaborative learning of personalized models over networks,” in *AISTATS*, 2017.
- [36] X. Lian, C. Zhang, H. Zhang, C.-J. Hsieh, W. Zhang, and J. Liu, “Can decentralized algorithms outperform centralized algorithms? A case study for decentralized parallel stochastic gradient descent,” in *NIPS*, 2017.
- [37] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, “Randomized gossip algorithms,” *IEEE Trans. Inform. Theor.*, vol. 52, no. 6, pp. 2508–2530, 2006.

- [38] M. Assran, N. Loizou, N. Ballas, and M. Rabbat, “Stochastic gradient push for distributed deep learning,” in *ICML*, 2019.
- [39] A. Bellet, R. Guerraoui, M. Taziki, and M. Tommasi, “Personalized and Private Peer-to-Peer Machine Learning,” in *AISTATS*, 2018.
- [40] I. Colin, A. Bellet, J. Salmon, and S. Cléménçon, “Gossip dual averaging for decentralized optimization of pairwise functions,” in *ICML*, 2016.
- [41] A. Elgabli, J. Park, A. S. Bedi, M. Bennis, and V. Aggarwal, “GADMM: Fast and communication efficient framework for distributed machine learning,” *arXiv preprint arXiv:1909.00047*, 2019.
- [42] A. Koloskova, S. U. Stich, and M. Jaggi, “Decentralized stochastic optimization and Gossip algorithms with compressed communication,” in *ICML*, 2019.
- [43] A. Lalitha, O. C. Kilinc, T. Javidi, and F. Koushanfar, “Peer-to-peer federated learning on graphs,” arXiv:1901.11173, Tech. Rep., 2019.
- [44] H. Tang, X. Lian, M. Yan, C. Zhang, and J. Liu, “D2: Decentralized training over decentralized data,” in *ICML*, 2018.
- [45] C. He, C. Tan, H. Tang, S. Qiu, and J. Liu, “Central server free federated learning over single-sided trust social networks,” *arXiv preprint arXiv:1910.04956*, 2019.
- [46] L. He, A. Bian, and M. Jaggi, “COLA: Decentralized linear learning,” in *NeurIPS 2018 – Advances in Neural Information Processing Systems*, vol. 31, pp. 4541–4551, 2018.
- [47] C. Yu, H. Tang, C. Renggli, S. Kassing, A. Singla, D. Alistarh, C. Zhang, and J. Liu, “Distributed learning over unreliable networks,” *arXiv preprint arXiv:1810.07766*, 2018.
- [48] G. Neglia, C. Xu, D. Towsley, and G. Calbi, “Decentralized gradient methods: Does topology matter?” In *AISTATS*, 2020.
- [49] J. Wang, A. Sahu, G. Joshi, and S. Kar, “MATCHA: Speeding up decentralized SGD via matching decomposition sampling,” *Preprint*, May 2019. [Online]. Available: <https://arxiv.org/abs/1905.09435>.

- [50] X. Lian, W. Zhang, C. Zhang, and J. Liu, “Asynchronous decentralized parallel stochastic gradient descent,” in *ICML*, 2018.
- [51] A. Koloskova, T. Lin, S. U. Stich, and M. Jaggi, “Decentralized deep learning with arbitrary communication compression,” *International Conference on Learning Representations (ICLR)*, 2020.
- [52] A. Koloskova, N. Loizou, S. Boreiri, M. Jaggi, and S. U. Stich, “A unified theory of decentralized SGD with changing topology and local updates,” in *ICML*, 2020.
- [53] J. Wang and G. Joshi, “Cooperative SGD: A unified framework for the design and analysis of communication-efficient SGD algorithms,” *Preprint*, Aug., 2018. [Online]. Available: <https://arxiv.org/abs/1808.07576>.
- [54] V. Zantedeschi, A. Bellet, and M. Tommasi, *Fully decentralized joint learning of personalized models and collaboration graphs*, Tech. Rep., arXiv:1901.08460, 2019.
- [55] A. Reisizadeh, H. Taheri, A. Mokhtari, H. Hassani, and R. Pedarsani, “Robust and communication-efficient collaborative learning,” *arXiv:1907.10595*, 2019.
- [56] H. Tang, X. Lian, S. Qiu, L. Yuan, C. Zhang, T. Zhang, and J. Liu, “DeepSqueeze: Parallel stochastic gradient descent with double-pass error-compensated compression,” *arXiv preprint arXiv:1907.07346*, 2019.
- [57] Z. Huang, S. Mitra, and N. Vaidya, “Differentially private distributed optimization,” in *ICDCN*, 2015.
- [58] E. Cyffers and A. Bellet, “Privacy amplification by decentralization,” *arXiv preprint arXiv:2012.05326*, 2020.
- [59] G. Wood, “Ethereum: A secure decentralised generalised transaction ledger,” *Ethereum Project Yellow Paper*, vol. 151, no. 2014, pp. 1–32, 2014.
- [60] K. A. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, “Practical secure aggregation for privacy-preserving machine learning,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, ACM, pp. 1175–1191, 2017.

- [61] R. Cheng, F. Zhang, J. Kos, W. He, N. Hynes, N. Johnson, A. Juels, A. Miller, and D. Song, “Ekiden: A platform for confidentiality-preserving, trustworthy, and performant smart contracts,” in *2019 IEEE European Symposium on Security and Privacy (EuroSecP)*, pp. 185–200, IEEE, 2019.
- [62] H. B. McMahan, D. Ramage, K. Talwar, and L. Zhang, “Learning differentially private recurrent language models,” in *International Conference on Learning Representations (ICLR)*, 2018.
- [63] Q. Yang, Y. Liu, T. Chen, and Y. Tong, “Federated machine learning: Concept and applications,” *CoRR*, vol. abs/1902.04885, 2019. [Online]. Available: <http://arxiv.org/abs/1902.04885>.
- [64] K. Cheng, T. Fan, Y. Jin, Y. Liu, T. Chen, and Q. Yang, “SecureBoost: A lossless federated learning framework,” *CoRR*, vol. abs/1901.08755, 2019. [Online]. Available: <http://arxiv.org/abs/1901.08755>.
- [65] S. Hardy, W. Henecka, H. Ivey-Law, R. Nock, G. Patrini, G. Smith, and B. Thorne, “Private federated learning on vertically partitioned data via entity resolution and additively homomorphic encryption,” *arXiv preprint arXiv:1711.10677*, 2017.
- [66] Y. Liu, Y. Kang, X. Zhang, L. Li, Y. Cheng, T. Chen, M. Hong, and Q. Yang, “A communication efficient vertical federated learning framework,” *CoRR*, vol. abs/1912.11187, 2019. [Online]. Available: <http://arxiv.org/abs/1912.11187>.
- [67] Y. Liu, Y. Kang, C. Xing, T. Chen, and Q. Yang, “A secure federated transfer learning framework,” *IEEE Intelligent Systems*, vol. 35, no. 4, pp. 70–82, 2020.
- [68] Y. Hu, P. Liu, L. Kong, and D. Niu, “Learning privately over distributed features: An ADMM sharing approach,” *arXiv preprint arXiv:1907.07735*, 2019.
- [69] Y. Liu, Z. Yi, and T. Chen, “Backdoor attacks and defenses in feature-partitioned collaborative learning,” *arXiv preprint arXiv:2007.03608*, 2020.
- [70] S. J. Pan and Q. Yang, “A survey on transfer learning,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.

- [71] H. Ludwig, N. Baracaldo, G. Thomas, Y. Zhou, A. Anwar, S. Rajamoni, Y. Ong, J. Radhakrishnan, A. Verma, M. Sinn, M. Purcell, A. Rawat, T. Minh, N. Holohan, S. Chakraborty, S. Whitherspoon, D. Steuer, L. Wynter, H. Hassan, S. Laguna, M. Yurochkin, M. Agarwal, E. Chuba, and A. Abay, “IBM federated learning: An enterprise framework white paper V0.1,” *arXiv preprint arXiv:2007.10987*, 2020.
- [72] J. Kang, Z. Xiong, D. Niyato, S. Xie, and J. Zhang, “Incentive mechanism for reliable federated learning: A joint optimization approach to combining reputation and contract theory,” *IEEE Internet of Things Journal*, 2019.
- [73] J. Kang, Z. Xiong, D. Niyato, H. Yu, Y.-C. Liang, and D. I. Kim, “Incentive design for efficient federated learning in mobile networks: A contract theory approach,” in *IEEE VTS Asia Pacific Wireless Communications Symposium, APWCS 2019, Singapore, August 28–30, 2019*, pp. 1–5, 2019.
- [74] H. Yu, Z. Liu, Y. Liu, T. Chen, M. Cong, X. Weng, D. Niyato, and Q. Yang, “A sustainable incentive scheme for federated learning,” *IEEE Intelligent Systems*, vol. 35, no. 4, pp. 58–69, 2020.
- [75] L. Lyu, J. Yu, K. Nandakumar, Y. Li, X. Ma, J. Jin, H. Yu, and K. S. Ng, “Towards fair and privacy-preserving federated deep models,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 31, no. 11, pp. 2524–2541, 2020.
- [76] Y. Kim, J. Sun, H. Yu, and X. Jiang, “Federated tensor factorization for computational phenotyping,” in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Halifax, NS, Canada, August 13–17, 2017*, pp. 887–895, 2017. DOI: [10.1145/3097983.3098118](https://doi.org/10.1145/3097983.3098118).
- [77] J. Ma, Q. Zhang, J. Lou, J. Ho, L. Xiong, and X. Jiang, “Privacy-preserving tensor factorization for collaborative health data analysis,” in *ACM CIKM*, vol. 2, 2019.
- [78] O. Gupta and R. Raskar, “Distributed learning of deep neural network over multiple agents,” *Journal of Network and Computer Applications*, vol. 116, pp. 1–8, 2018.

- [79] P. Vepakomma, O. Gupta, T. Swedish, and R. Raskar, “Split learning for health: Distributed deep learning without sharing raw patient data,” *arXiv preprint arXiv:1812.00564*, 2018.
- [80] I. Ceballos, V. Sharma, E. Mugica, A. Singh, A. Roman, P. Vepakomma, and R. Raskar, “SplitNN-driven vertical partitioning,” *arXiv preprint arXiv:2008.04137*, 2018.
- [81] A. Singh, P. Vepakomma, O. Gupta, and R. Raskar, “Detailed comparison of communication efficiency of split learning and federated learning,” *arXiv preprint arXiv:1909.09145*, 2019.
- [82] Z. Huo, B. Gu, and H. Huang, “Training neural networks using features replay,” in *Advances in Neural Information Processing Systems*, pp. 6659–6668, 2018.
- [83] M. Jaderberg, W. M. Czarnecki, S. Osindero, O. Vinyals, A. Graves, D. Silver, and K. Kavukcuoglu, “Decoupled neural interfaces using synthetic gradients,” in *Proceedings of the 34th International Conference on Machine Learning – Volume 70*, JMLR.org, pp. 1627–1635, 2017.
- [84] V. Sharma, P. Vepakomma, T. Swedish, K. Chang, J. Kalpathy-Cramer, and R. Raskar, “ExpertMatcher: Automating ML model selection for clients using hidden representations,” *arXiv preprint arXiv:1910.03731*, 2019.
- [85] P. Vepakomma, O. Singh, A. Gupta, and R. Raskar, “Nopeek: Information leakage reduction to share activations in distributed deep learning,” *arXiv preprint arXiv:2008.09161*, 2020.
- [86] G. J. Székely, M. L. Rizzo, and N. K. Bakirov, “Measuring and testing dependence by correlation of distances,” *The Annals of Statistics*, vol. 35, no. 6, pp. 2769–2794, 2007.
- [87] P. Vepakomma, C. Tonde, and A. Elgammal, “Supervised dimensionality reduction via distance correlation maximization,” *Electronic Journal of Statistics*, vol. 12, no. 1, pp. 960–984, 2018.
- [88] A. Singh, A. Chopra, V. Sharma, E. Garza, E. Zhang, P. Vepakomma, and R. Raskar, “DISCO: Dynamic and invariant sensitive channel obfuscation for deep neural networks,” *arXiv preprint arXiv:2012.11025*, 2020.

- [89] J. G. Moreno-Torres, T. Raeder, R. Alaiz-Rodríguez, N. V. Chawla, and F. Herrera, “A unifying view on dataset shift in classification,” *Pattern Recogn.*, vol. 45, no. 1, Jan. 2012.
- [90] J. Quionero-Candela, M. Sugiyama, A. Schwaighofer, and N. D. Lawrence, *Dataset Shift in Machine Learning*. The MIT Press, 2009.
- [91] K. Hsieh, A. Phanishayee, O. Mutlu, and P. B. Gibbons, *The non-IID data quagmire of decentralized machine learning*, 2019. [Online]. Available: <https://arxiv.org/abs/1910.00189>.
- [92] H. Eichner, T. Koren, H. B. McMahan, N. Srebro, and K. Talwar, “Semi-cyclic stochastic gradient descent,” in *Accepted to ICML 2019*, 2019. [Online]. Available: <https://arxiv.org/abs/1904.10120>.
- [93] T. Wang, J.-Y. Zhu, A. Torralba, and A. A. Efros, “Dataset distillation,” *arXiv preprint arXiv:1811.10959*, 2018.
- [94] B. Woodworth, J. Wang, H. B. McMahan, and N. Srebro, “Graph oracle models, lower bounds, and gaps for parallel stochastic optimization,” in *Advances in Neural Information Processing Systems (NIPS)*, 2018. [Online]. Available: <https://arxiv.org/abs/1805.10222>.
- [95] A. Cotter, O. Shamir, N. Srebro, and K. Sridharan, “Better mini-batch algorithms via accelerated gradient methods,” in *Advances in Neural Information Processing Systems*, J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Q. Weinberger, Eds., vol. 24, pp. 1647–1655, Curran Associates, Inc., 2011. [Online]. Available: <https://proceedings.neurips.cc/paper/2011/file/b55ec28c52d5f6205684a473a2193564-Paper.pdf>.
- [96] O. Dekel, R. Gilad-Bachrach, O. Shamir, and L. Xiao, “Optimal distributed online prediction using mini-batches,” *J. Mach. Learn. Res.*, vol. 13, no. 6, pp. 165–202, 2012.
- [97] G. Lan, “An optimal method for stochastic composite optimization,” *Math. Program.*, vol. 133, no. 1, pp. 365–397, 2012.
- [98] H. Yu, S. Yang, and S. Zhu, “Parallel restarted SGD for non-convex optimization with faster convergence and less communication,” *arXiv preprint arXiv:1807.06629*, 2018.

- [99] S. U. Stich, “Local SGD converges fast and communicates little,” in *International Conference on Learning Representations (ICLR)*, 2019.
- [100] S. U. Stich and S. P. Karimireddy, “The error-feedback framework: Better rates for SGD with delayed gradients and compressed communication,” *arXiv:1909.05350*, 2019.
- [101] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh, “Scaffold: Stochastic controlled averaging for federated learning,” in *International Conference on Machine Learning*, PMLR, pp. 5132–5143, 2020.
- [102] A. Khaled, K. Mishchenko, and P. Richtárik, *Better communication complexity for local SGD*, 2019. [Online]. Available: <https://arxiv.org/abs/1909.04746>.
- [103] T. Lin, S. U. Stich, and M. Jaggi, “Don’t use large mini-batches, use local SGD,” *International Conference on Learning Representations (ICLR)*, 2020.
- [104] K. K. Patel and A. Dieuleveut, “Communication trade-offs for synchronized distributed SGD with large step size,” *NeurIPS*, 2019.
- [105] A. Reisizadeh, A. Mokhtari, H. Hassani, A. Jadbabaie, and R. Pedarsani, “Fedpaq: A communication-efficient federated learning method with periodic averaging and quantization,” *arXiv preprint arXiv:1909.13014*, 2019.
- [106] B. Woodworth, K. K. Patel, S. U. Stich, Z. Dai, B. Bullins, H. B. McMahan, O. Shamir, and N. Srebro, “Is local SGD better than minibatch SGD?” *arXiv preprint arXiv:2002.07839*, 2020.
- [107] S. Zhang, A. E. Choromanska, and Y. LeCun, “Deep learning with elastic averaging SGD,” in *Advances in Neural Information Processing Systems*, pp. 685–693, 2015.
- [108] F. Haddadpour, M. M. Kamani, M. Mahdavi, and V. R. Cadambe, “Local SGD with periodic averaging: Tighter analysis and adaptive synchronization,” *arXiv preprint arXiv:1910.13598*, 2019.

- [109] H. Karimi, J. Nutini, and M. Schmidt, “Linear convergence of gradient and proximal-gradient methods under the Polyak-Łojasiewicz condition,” in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, Springer, pp. 795–811, 2016.
- [110] J. Wang and G. Joshi, “Adaptive communication strategies for best error-runtime trade-offs in communication-efficient distributed SGD,” in *Proceedings of the SysML Conference*, Apr. 2019. [Online]. Available: <https://arxiv.org/abs/1810.08313>.
- [111] J. Wang, Q. Liu, H. Liang, G. Joshi, and H. V. Poor, “Tackling the objective inconsistency problem in heterogeneous federated optimization,” in *Advances in Neural Information Processing Systems*, vol. 33, pp. 7611–7623, 2020.
- [112] S. Dutta, G. Joshi, S. Ghosh, P. Dube, and P. Nagpurkar, “Slow and stale gradients can win the race: Error-runtime trade-offs in distributed SGD,” *International Conference on Artificial Intelligence and Statistics (AISTATS)*, Apr. 2018. [Online]. Available: <https://arxiv.org/abs/1803.01113>.
- [113] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, *Federated optimization in heterogeneous networks*, 2018. [Online]. Available: <https://arxiv.org/abs/1812.06127>.
- [114] X. Li, W. Yang, S. Wang, and Z. Zhang, “Communication efficient decentralized training with multiple local updates,” *arXiv preprint arXiv:1910.09126*, 2019.
- [115] A. Khaled, K. Mishchenko, and P. Richtárik, *First analysis of local GD on heterogeneous data*, 2019. [Online]. Available: <https://arxiv.org/abs/1909.04715>.
- [116] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, “On the convergence of FedAvg on non-IID data,” *arXiv preprint arXiv:1907.02189*, 2019.
- [117] H. Yu, R. Jin, and S. Yang, “On the linear speedup analysis of communication efficient momentum SGD for distributed non-convex optimization,” *arXiv preprint arXiv:1905.03817*, 2019.
- [118] C. Xie, O. Koyejo, I. Gupta, and H. Lin, “Local adaalter: Communication-efficient stochastic gradient descent with adaptive learning rates,” *arXiv preprint arXiv:1911.09030*, 2019.

- [119] J. Duchi, E. Hazan, and Y. Singer, “Adaptive subgradient methods for online learning and stochastic optimization,” *J. Mach. Learn. Res.*, vol. 12, no. 61, pp. 2121–2159, 2011.
- [120] H. B. McMahan and M. Streeter, “Adaptive bound optimization for online convex optimization,” *arXiv preprint arXiv:1002.4908*, 2010.
- [121] S. Reddi, Z. Charles, M. Zaheer, Z. Garrett, K. Rush, J. Konečný, S. Kumar, and H. B. McMahan, “Adaptive federated optimization,” *arXiv preprint arXiv:2003.00295*, 2020.
- [122] T.-M. H. Hsu, H. Qi, and M. Brown, “Measuring the effects of non-identical data distribution for federated visual classification,” *arXiv preprint arXiv:1909.06335*, 2019.
- [123] J. Wang, V. Tantia, N. Ballas, and M. Rabbat, “SlowMo: Improving communication-efficient distributed SGD with slow momentum,” *arXiv preprint arXiv:1910.00643*, 2019.
- [124] S. P. Karimireddy, M. Jaggi, S. Kale, M. Mohri, S. J. Reddi, S. U. Stich, and A. T. Suresh, “Mime: Mimicking centralized stochastic algorithms in federated learning,” *arXiv preprint arXiv:2008.03606*, 2020.
- [125] I. Almeida and J. Xavier, “DJAM: Distributed Jacobi asynchronous method for learning personal models,” *IEEE Signal Process. Lett.*, vol. 25, no. 9, pp. 1389–1392, 2018.
- [126] K. Wang, R. Mathews, C. Kiddon, H. Eichner, F. Beaufays, and D. Ramage, “Federated evaluation of on-device personalization,” *arXiv preprint arXiv:1910.10252*, 2019.
- [127] Y. Zhang and Q. Yang, “A survey on multi-task learning,” *CoRR*, vol. abs/1707.08114, 2017. [Online]. Available: <http://arxiv.org/abs/1707.08114>.
- [128] V. Smith, C.-K. Chiang, M. Sanjabi, and A. S. Talwalkar, “Federated multi-task learning,” in *NIPS*, 2017.
- [129] Y. Mansour, M. Mohri, J. Ro, and A. T. Suresh, “Three approaches for personalization with applications to federated learning,” *arXiv preprint arXiv:2002.10619*, 2020.
- [130] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. W. Vaughan, “A theory of learning from different domains,” *Mach. Learn.*, vol. 79, no. 1–2, pp. 151–175, 2010.

- [131] C. Cortes and M. Mohri, “Domain adaptation and sample bias correction theory and algorithm for regression,” *Theor. Comput. Sci.*, vol. 519, pp. 103–126, 2014.
- [132] C. Cortes, M. Mohri, A. T. Suresh, and N. Zhang, “Multiple-source adaptation with domain classifiers,” *arXiv preprint arXiv:2008.11036*, 2020.
- [133] Y. Mansour, M. Mohri, and A. Rostamizadeh, “Domain adaptation: Learning bounds and algorithms,” *arXiv preprint arXiv:0902.3430*, 2009.
- [134] Y. Mansour, M. Mohri, A. T. Suresh, and K. Wu, “A theory of multiple-source adaptation with limited target labeled data,” *arXiv preprint arXiv:2007.09762*, 2020.
- [135] J. Baxter, “A model of inductive bias learning,” *J. Artif. Intell. Res.*, vol. 12, pp. 149–198, 2000.
- [136] C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *Proceedings of the 34th International Conference on Machine Learning*, 2017.
- [137] A. Nichol, J. Achiam, and J. Schulman, “On first-order meta-learning algorithms,” *arXiv preprint arXiv:1803.02999*, 2018.
- [138] M. Khodak, M.-F. Balcan, and A. Talwalkar, “Adaptive gradient-based meta-learning methods,” in *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [139] Y. Jiang, J. Konečný, K. Rush, and S. Kannan, “Improving federated learning personalization via model agnostic meta learning,” *arXiv preprint arXiv:1909.12488*, 2019.
- [140] K. C. Sim, F. Beaufays, A. Benard, D. Guliani, A. Kabel, N. Khare, T. Lucassen, P. Zadrazil, H. Zhang, L. Johnson, G. Motta, and L. Zhou, “Personalization of end-to-end speech recognition on mobile devices for named entities,” *arXiv preprint arXiv:1912.09251*, 2019.
- [141] A. Fallah, A. Mokhtari, and A. Ozdaglar, “Personalized federated learning: A meta-learning approach,” *arXiv preprint arXiv:2002.07948*, 2020.
- [142] J. Li, M. Khodak, S. Caldas, and A. Talwalkar, “Differentially private meta-learning,” *arXiv preprint arXiv:1909.05830*, 2019.

- [143] B. M. Lake, R. Salakhutdinov, J. Gross, and J. B. Tenenbaum, “One shot learning of simple visual concepts,” in *Proceedings of the Conference of the Cognitive Science Society (CogSci)*, 2017.
- [144] S. Ravi and H. Larochelle, “Optimization as a model for few-shot learning,” in *Proceedings of the 5th International Conference on Learning Representations*, 2017.
- [145] D. L. Silver, Q. Yang, and L. Li, “Lifelong machine learning systems: Beyond learning algorithms,” in *AAAI Spring Symposium Series*, 2013.
- [146] J. Snell, K. Swersky, and R. S. Zemel, “Prototypical networks for few-shot learning,” in *Advances in Neural Information Processing Systems*, pp. 4080–4090, 2017.
- [147] J. Hoffman, M. Mohri, and N. Zhang, “Algorithms and theory for multiple-source adaptation,” in *Advances in Neural Information Processing Systems*, pp. 8246–8256, 2018.
- [148] Y. Mansour, M. Mohri, and A. Rostamizadeh, “Domain adaptation with multiple sources,” in *Advances in Neural Information Processing Systems*, pp. 1041–1048, 2009.
- [149] M. Mohri, G. Sivek, and A. T. Suresh, “Agnostic Federated Learning,” in *ICML*, 2019.
- [150] P. Christen, *Data Matching: Concepts and Techniques for Record Linkage, Entity Resolution, and Duplicate Detection*. Springer Science & Business Media, 2012.
- [151] R. Schnell, T. Bachteler, and J. Reiher, “A novel error-tolerant anonymous linking code,”
- [152] R. Schnell, “Efficient private record linkage of very large datasets,” in *59th World Statistics Congress*, 2013.
- [153] T. Yu, E. Bagdasaryan, and V. Shmatikov, “Salvaging federated learning by local adaptation,” *arXiv preprint arXiv:2002.04758*, 2020.
- [154] G. Patrini, R. Nock, S. Hardy, and T. S. Caetano, “Fast learning from distributed datasets without entity matching,” in *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9–15 July 2016*, pp. 1909–1917, 2016. [Online]. Available: <http://www.ijcai.org/Abstract/16/273>.

- [155] R. D. King, C. Feng, and A. Sutherland, “StatLog: Comparison of classification algorithms on large real-world problems,” *Appl. Artif. Intell.*, vol. 9, no. 3, pp. 289–333, 1995.
- [156] R. Kohavi and G. H. John, “Automatic parameter selection by minimizing estimated error,” in *Machine Learning Proceedings 1995*, Elsevier, 1995, pp. 304–312.
- [157] B. D. Ripley, “Statistical aspects of neural networks,” *Networks and Chaos—Statistical and Probabilistic Aspects*, vol. 50, Chapman and Hall, 1993, pp. 40–123.
- [158] J. S. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, “Algorithms for hyper-parameter optimization,” in *Advances in Neural Information Processing Systems*, pp. 2546–2554, 2011.
- [159] S. Falkner, A. Klein, and F. Hutter, “BOHB: Robust and efficient hyperparameter optimization at scale,” *arXiv preprint arXiv:1807.01774*, 2018.
- [160] F. Pedregosa, “Hyperparameter optimization with approximate gradient,” *arXiv preprint arXiv:1602.02355*, 2016.
- [161] J. Snoek, O. Rippel, K. Swersky, R. Kiros, N. Satish, N. Sundaram, M. Patwary, M. Prabhat, and R. Adams, “Scalable Bayesian optimization using deep neural networks,” in *International Conference on Machine Learning*, pp. 2171–2180, 2015.
- [162] Z. Charles and J. Konečný, “On the outsized importance of learning rates in local update methods,” *arXiv preprint arXiv:2007.00878*, 2020.
- [163] K. A. Bonawitz, F. Salehi, J. Konečný, B. McMahan, and M. Gruteser, “Federated learning with autotuned communication-efficient secure aggregation,” in *2019 53rd Asilomar Conference on Signals, Systems, and Computers*, IEEE, 2019.
- [164] O. Thakkar, G. Andrew, and H. B. McMahan, “Differentially private learning with adaptive clipping,” *arXiv preprint arXiv:1905.03871*, 2019.
- [165] I. Bello, B. Zoph, V. Vasudevan, and Q. V. Le, “Neural optimizer search with reinforcement learning,” in *Proceedings of the 34th International Conference on Machine Learning – Volume 70*, JMLR.org, pp. 459–468, 2017.

- [166] T. Elsken, J. Hendrik Metzen, and F. Hutter, “Efficient multi-objective neural architecture search via Lamarckian evolution,” *arXiv preprint arXiv:1804.09081*, 2018.
- [167] H. Liu, K. Simonyan, and Y. Yang, “DARTS: Differentiable architecture search,” *arXiv preprint arXiv:1806.09055*, 2018.
- [168] R. Luo, F. Tian, T. Qin, E. Chen, and T.-Y. Liu, “Neural architecture optimization,” in *Advances in Neural Information Processing Systems*, pp. 7816–7827, 2018.
- [169] C. He, H. Ye, L. Shen, and T. Zhang, “Milenas: Efficient neural architecture search via mixed-level reformulation,” in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [170] E. Real, S. Moore, A. Selle, S. Saxena, Y. L. Suematsu, J. Tan, Q. V. Le, and A. Kurakin, “Large-scale evolution of image classifiers,” in *Proceedings of the 34th International Conference on Machine Learning – Volume 70*, JMLR.org, pp. 2902–2911, 2017.
- [171] E. Real, A. Aggarwal, Y. Huang, and Q. V. Le, “Regularized evolution for image classifier architecture search,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, pp. 4780–4789, 2019.
- [172] H. Pham, M. Guan, B. Zoph, Q. Le, and J. Dean, “Efficient neural architecture search via parameter sharing,” in *International Conference on Machine Learning*, pp. 4092–4101, 2018.
- [173] S. Xie, H. Zheng, C. Liu, and L. Lin, “SNAS: Stochastic neural architecture search,” *arXiv preprint arXiv:1812.09926*, 2018.
- [174] A. Gaier and D. Ha, “Weight agnostic neural networks,” *arXiv preprint arXiv:1906.04358*, 2019.
- [175] C. He, M. Annavaram, and S. Avestimehr, “FedNAS: Federated deep learning via neural architecture search,” *arXiv preprint arXiv:2004.08546*, 2020.
- [176] S. Augenstein, H. B. McMahan, D. Ramage, S. Ramaswamy, P. Kairouz, M. Chen, R. Mathews, and B. A. y Arcas, *Generative models for effective ML on private, decentralized datasets*, 2019. [Online]. Available: <https://arxiv.org/abs/1911.06679>.

- [177] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, “Federated learning: Strategies for improving communication efficiency,” *arXiv preprint arXiv:1610.05492*, 2016.
- [178] J. Acharya, C. L. Canonne, and H. Tyagi, “Inference under information constraints I: Lower bounds from chi-square contraction,” *IEEE Trans. Inform. Theor.*, vol. 66, no. 12, pp. 7835–7855, 2020.
- [179] L. P. Barnes, Y. Han, and A. Ozgur, “Lower bounds for learning distributions under communication constraints via Fisher information,” *J. Mach. Learn. Res.*, vol. 21, no. 236, pp. 1–30, 2020.
- [180] M. Braverman, A. Garg, T. Ma, H. L. Nguyen, and D. P. Woodruff, “Communication lower bounds for statistical estimation problems via a distributed data processing inequality,” in *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, pp. 1011–1020, ACM, 2016.
- [181] L. P. Barnes, H. A. Inan, B. Isik, and A. Ozgur, “rTop-k: A statistical estimation approach to distributed SGD,” *arXiv preprint arXiv:2005.10761*, 2020.
- [182] Y. Zhang, J. Duchi, M. I. Jordan, and M. J. Wainwright, “Information-theoretic lower bounds for distributed statistical estimation with communication constraints,” in *Advances in Neural Information Processing Systems*, pp. 2328–2336, 2013.
- [183] Y. Han, A. Özgür, and T. Weissman, “Geometric lower bounds for distributed parameter estimation under communication constraints,” in *Proceedings of Machine Learning Research*, pp. 1–26, 75, 2018.
- [184] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic, “QSGD: Communication-efficient SGD via gradient quantization and encoding,” in *NIPS – Advances in Neural Information Processing Systems*, pp. 1709–1720, 2017.
- [185] D. Basu, D. Data, C. Karakus, and S. N. Diggavi, “Qsparse-local-SGD: Distributed SGD with quantization, sparsification, and local computations,” *IEEE Sel. Areas Inf. Theor.*, vol. 1, no. 1, pp. 217–226, 2020.

- [186] S. Horvath, C.-Y. Ho, L. Horvath, A. N. Sahu, M. Canini, and P. Richtárik, “Natural compression for distributed deep learning,” *arXiv preprint arXiv:1905.10988*, 2019.
- [187] J. Konečný and P. Richtárik, “Randomized distributed mean estimation: Accuracy vs. communication,” *Frontiers in Applied Mathematics and Statistics*, vol. 4, p. 62, 2018.
- [188] A. T. Suresh, F. X. Yu, S. Kumar, and H. B. McMahan, “Distributed mean estimation with limited communication,” in *Proceedings of the 34th International Conference on Machine Learning – Volume 70*, JMLR.org, pp. 3329–3337, 2017.
- [189] S. Chraïbi, A. Khaled, D. Kovalev, P. Richtárik, A. Salim, and M. Takáč, “Distributed fixed point methods with compressed iterates,” *arXiv preprint arXiv:1912.09925*, 2019.
- [190] S. Caldas, J. Konečný, H. B. McMahan, and A. Talwalkar, “Expanding the reach of federated learning by reducing client resource requirements,” *arXiv preprint arXiv:1812.07210*, 2018.
- [191] J. Hamer, M. Mohri, and A. T. Suresh, “Fedboost: A communication-efficient algorithm for federated learning,” in *International Conference on Machine Learning*, PMLR, pp. 3973–3983, 2020.
- [192] C. He, M. Annavaram, and S. Avestimehr, “Group knowledge transfer: Federated learning of large cnns at the edge,” in *Advances in Neural Information Processing Systems*, vol. 34, pp. 14068–14080, 2020.
- [193] S. P. Karimireddy, Q. Rebjock, S. Stich, and M. Jaggi, “Error feedback fixes SignSGD and other gradient compression schemes,” in *ICML*, 2019.
- [194] Y. Lin, S. Han, H. Mao, Y. Wang, and W. J. Dally, “Deep gradient compression: Reducing the communication bandwidth for distributed training,” *arXiv preprint arXiv:1712.01887*, 2017.
- [195] F. Sattler, S. Wiedemann, K.-R. Müller, and W. Samek, “Robust and communication-efficient federated learning from non-IID data,” *arXiv preprint arXiv:1903.02891*, 2019.

- [196] T. Vogels, S. P. Karimireddy, and M. Jaggi, “PowerSGD: Practical low-rank gradient compression for distributed optimization,” in *NeurIPS 2019 – Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [197] D. Blalock, J. J. G. Ortiz, J. Frankle, and J. Gutttag, “What is the state of neural network pruning?” *arXiv preprint arXiv:2003.03033*, 2020.
- [198] M. Courbariaux, Y. Bengio, and J.-P. David, “BinaryConnect: Training deep neural networks with binary weights during propagations,” in *Advances in Neural Information Processing Systems*, pp. 3123–3131, 2015.
- [199] S. Han, H. Mao, and W. J. Dally, “Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding,” *arXiv preprint arXiv:1510.00149*, 2015.
- [200] D. Lin, S. Talathi, and S. Annapureddy, “Fixed point quantization of deep convolutional networks,” in *International Conference on Machine Learning*, pp. 2849–2858, 2016.
- [201] D. Oktay, J. Ballé, S. Singh, and A. Shrivastava, “Model compression by entropy penalized reparameterization,” *arXiv preprint arXiv:1906.06624*, 2019.
- [202] M. Zhu and S. Gupta, “To prune, or not to prune: Exploring the efficacy of pruning for model compression,” *arXiv preprint arXiv:1710.01878*, 2017.
- [203] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley & Sons, 2012.
- [204] X. Wu, R. Guo, A. T. Suresh, S. Kumar, D. N. Holtmann-Rice, D. Simcha, and F. X. Yu, “Multiscale quantization for fast similarity search,” in *Advances in Neural Information Processing Systems*, pp. 5745–5755, 2017.
- [205] V. Gandikota, R. K. Maity, and A. Mazumdar, “VqSGD: Vector quantized stochastic gradient descent,” *arXiv preprint arXiv:1911.07971*, 2019.
- [206] K. A. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, “Practical secure aggregation for federated learning on user-held data,” *arXiv preprint arXiv:1611.04482*, 2016.

- [207] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, “Deep learning with differential privacy,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, ACM, pp. 308–318, 2016.
- [208] J. H. Bell, K. A. Bonawitz, A. Gascón, T. Lepoint, and M. Raykova, “Secure single-server aggregation with (poly)logarithmic overhead,” in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, ACM, pp. 1253–1269, 2020.
- [209] N. Agarwal, A. T. Suresh, F. X. Yu, S. Kumar, and B. McMahan, “CpSGD: Communication-efficient and differentially-private distributed SGD,” in *Advances in Neural Information Processing Systems*, pp. 7564–7575, 2018.
- [210] E. Jeong, S. Oh, H. Kim, J. Park, M. Bennis, and S.-L. Kim, “Communication-efficient on-device machine learning: Federated distillation and augmentation under non-IID private data,” *CoRR*, vol. abs/1811.11479, 2018. [Online]. Available: <http://arxiv.org/abs/1811.11479>.
- [211] J. Park, S. Samarakoon, M. Bennis, and M. Debbah, “Wireless network intelligence at the edge,” *CoRR*, vol. abs/1812.02858, 2018. [Online]. Available: <http://arxiv.org/abs/1812.02858>.
- [212] S. Samarakoon, M. Bennis, W. Saad, and M. Debbah, “Federated learning for ultra-reliable low-latency V2V communications,” *CoRR*, vol. abs/1805.09253, 2018. [Online]. Available: <http://arxiv.org/abs/1805.09253>.
- [213] O. Abari, H. Rahul, and D. Katabi, “Over-the-air function computation in sensor networks,” *CoRR*, vol. abs/1612.02307, 2016. [Online]. Available: <http://arxiv.org/abs/1612.02307>.
- [214] A. Elgabli, J. Park, C. B. Issaid, and M. Bennis, “Harnessing wireless channels for scalable and privacy-preserving federated learning,” *IEEE Trans. Comm.*, to be published.
- [215] Y. Zhao, C. Yu, P. Zhao, and J. Liu, “Decentralized online learning: Take benefits from others’ data without sharing your own to track global trend,” *arXiv preprint arXiv:1901.10593*, 2019.

- [216] K. Shridhar, F. Laumann, and M. Liwicki, “A comprehensive guide to Bayesian convolutional neural network with variational inference,” *arXiv preprint arXiv:1901.02731*, 2019.
- [217] A. Lalitha, X. Wang, O. Kilinc, Y. Lu, T. Javidi, and F. Koushanfar, “Decentralized Bayesian learning over graphs,” *arXiv preprint: arXiv:1905.10466*, 2019.
- [218] F. Mireshghallah, M. Taram, P. Vepakomma, A. Singh, R. Raskar, and E. Hadi, “Privacy in deep learning: A survey,” *arXiv preprint arXiv:2004.12254*, 2020.
- [219] A. Bittau, Ú. Erlingsson, P. Maniatis, I. Mironov, A. Raghunathan, D. Lie, M. Rudominer, U. Kode, J. Tinnes, and B. Seefeld, “Prochlo: Strong privacy for analytics in the crowd,” in *Proceedings of the 26th Symposium on Operating Systems Principles, SOSP '17*, pp. 441–459, New York, NY, USA: ACM, 2017. DOI: [10.1145/3132747.3132769](https://doi.org/10.1145/3132747.3132769).
- [220] NSA, *Defense in depth: A practical strategy for achieving Information Assurance in today's highly networked environments*, 2012. [Online]. Available: <https://apps.nsa.gov/iaarchive/library/ia-guidance/archive/defense-in-depth.cfm>.
- [221] C.-C. Y. Andrew, “How to generate and exchange secrets (extended abstract),” in *FOCS*, pp. 162–167, IEEE Computer Society, 1986.
- [222] T. Araki, J. Furukawa, Y. Lindell, A. Nof, and K. Ohara, “High-throughput semi-honest secure three-party computation with an honest majority,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, ACM, pp. 805–817, 2016.
- [223] P. Bogetoft, D. L. Christensen, I. Damgård, M. Geisler, T. P. Jakobsen, M. Krøigaard, J. D. Nielsen, J. B. Nielsen, K. Nielsen, J. Pagter, M. I. Schwartzbach, and T. Toft, “Secure multiparty computation goes live,” in *Financial Cryptography*, Lecture Notes in Computer Science, vol. 5628, pp. 325–343, Springer, 2009.
- [224] D. Bogdanov, R. Talviste, and J. Willemsen, “Deploying secure multi-party computation for financial data analysis – (short paper),” in *Financial Cryptography*, Lecture Notes in Computer Science, vol. 7397, pp. 57–64, Springer, 2012.

- [225] A. Lapets, N. Volgushev, A. Bestavros, F. Jansen, and M. Varia, “Secure MPC for analytics as a web application,” in *SecDev*, pp. 73–74, IEEE Computer Society, 2016.
- [226] J. Furukawa, Y. Lindell, A. Nof, and O. Weinstein, “High-throughput secure three-party computation for malicious adversaries and an honest majority,” in *EUROCRYPT (2)*, Lecture Notes in Computer Science, vol. 10211, pp. 225–255, 2017.
- [227] M. Ion, B. Kreuter, E. Nergiz, S. Patel, S. Saxena, K. Seth, D. Shanahan, and M. Yung, “Private intersection-sum protocol with applications to attributing aggregate ad conversions,” *IACR Cryptology ePrint Archive*, vol. 2017, p. 738, 2017.
- [228] M. Ion, B. Kreuter, A. E. Nergiz, S. Patel, M. Raykova, S. Saxena, K. Seth, D. Shanahan, and M. Yung, “On deploying secure computing commercially: Private intersection-sum protocols and their business applications,” *IACR Cryptology ePrint Archive*, vol. 2019, p. 723, 2019.
- [229] Y. Ishai, J. Kilian, K. Nissim, and E. Petrank, “Extending oblivious transfers efficiently,” in *CRYPTO*, Lecture Notes in Computer Science, vol. 2729, pp. 145–161, Springer, 2003.
- [230] F. Bourse, M. Minelli, M. Minihold, and P. Paillier, “Fast homomorphic evaluation of deep discretized neural networks,” in *CRYPTO (3)*, Lecture Notes in Computer Science, vol. 10993, pp. 483–512, Springer, 2018.
- [231] R. Gilad-Bachrach, N. Dowlin, K. Laine, K. E. Lauter, M. Naehrig, and J. Wernsing, “CryptoNets: Applying neural networks to encrypted data with high throughput and accuracy,” in *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19–24, 2016*, pp. 201–210, 2016. [Online]. Available: <http://proceedings.mlr.press/v48/gilad-bachrach16.html>.
- [232] A. Gascón, P. Schoppmann, B. Balle, M. Raykova, J. Doerner, S. Zahur, and D. Evans, “Privacy-preserving distributed linear regression on high-dimensional data,” *PoPETs*, vol. 2017, no. 4, pp. 345–364, 2017.

- [233] N. Agrawal, A. S. Shamsabadi, M. J. Kusner, and A. Gascón, “QUOTIENT: Two-party secure neural network training and prediction,” in *Proceedings of the ACM Conference on Computer and Communication Security (CCS)*, 2019.
- [234] O. Goldreich, S. Micali, and A. Wigderson, “How to play any mental game,” in *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, STOC '87, pp. 218–229, New York, New York, USA: ACM, 1987. DOI: [10.1145/28395.28420](https://doi.org/10.1145/28395.28420).
- [235] V. Nikolaenko, U. Weinsberg, S. Ioannidis, M. Joye, D. Boneh, and N. Taft, “Privacy-preserving ridge regression on hundreds of millions of records,” in *IEEE Symposium on Security and Privacy*, pp. 334–348, IEEE Computer Society, 2013.
- [236] P. Mohassel and Y. Zhang, “SecureML: A system for scalable privacy-preserving machine learning,” in *IEEE Symposium on Security and Privacy*, pp. 19–38, IEEE Computer Society, 2017.
- [237] A. Barak, D. Escudero, A. P. K. Dalskov, and M. Keller, “Secure evaluation of quantized neural networks,” *IACR Cryptology ePrint Archive*, vol. 2019, p. 131, 2019. [Online]. Available: <https://eprint.iacr.org/2019/131>.
- [238] C. Gentry, “Fully homomorphic encryption using ideal lattices,” in *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, pp. 169–178, 2009.
- [239] Z. Brakerski, “Fully homomorphic encryption without modulus switching from classical GapSVP,” in *CRYPTO*, Lecture Notes in Computer Science, vol. 7417, pp. 868–886, Springer, 2012.
- [240] J. Fan and F. Vercauteren, “Somewhat practical fully homomorphic encryption,” *IACR Cryptology ePrint Archive*, vol. 2012, p. 144, 2012.
- [241] Z. Brakerski, C. Gentry, and V. Vaikuntanathan, “(leveled) Fully homomorphic encryption without bootstrapping,” in *ITCS*, pp. 309–325, ACM, 2012.
- [242] J.-S. Coron, T. Lepoint, and M. Tibouchi, “Scale-invariant fully homomorphic encryption over the integers,” in *Public Key Cryptography*, Lecture Notes in Computer Science, vol. 8383, pp. 311–328, Springer, 2014.

- [243] *HElib*, <https://github.com/homenc/HElib>, Oct. 2019.
- [244] *Lattigo 2.0.0*, [Online]: <http://github.com/ldsec/lattigo>, Oct. 2020, EPFL-LDS.
- [245] *PALISADE lattice cryptography library*, <https://gitlab.com/palisade/palisade-release>, Oct. 2019.
- [246] *Microsoft SEAL (release 3.6)*, <https://github.com/Microsoft/SEAL>, Nov. 2020, Microsoft Research, Redmond, WA.
- [247] *SHELL*, <https://github.com/google/shell-encryption>, Dec. 2020. Google.
- [248] S. S. Sathya, P. Vepakomma, R. Raskar, R. Ramachandra, and S. Bhattacharya, “A review of homomorphic encryption libraries for secure computation,” *arXiv preprint arXiv:1812.02428*, 2018.
- [249] M. Chenal and Q. Tang, “On key recovery attacks against existing somewhat homomorphic encryption schemes,” in *LATINCRYPT*, Lecture Notes in Computer Science, vol. 8895, pp. 239–258, Springer, 2014.
- [250] L. Reyzin, A. D. Smith, and S. Yakoubov, “Turning HATE into LOVE: Homomorphic ad hoc threshold encryption for scalable MPC,” *IACR Cryptology ePrint Archive*, vol. 2018, p. 997, 2018.
- [251] E. Roth, D. Noble, B. H. Falk, and A. Haeberlen, “Honeycrisp: Large-scale differentially private aggregation without a trusted core,” in *SOSP*, pp. 196–210, ACM, 2019.
- [252] P. Subramanyan, R. Sinha, I. Lebedev, S. Devadas, and S. A. Seshia, “A formal foundation for secure remote execution of enclaves,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, ACM, pp. 2435–2450, 2017.
- [253] R. Intel, “Architecture instruction set extensions programming reference,” *Intel Corporation*, Feb., 2012.
- [254] V. Costan and S. Devadas, “Intel SGX explained,” *IACR Cryptology ePrint Archive*, vol. 2016, no. 086, pp. 1–118, 2016.
- [255] Arm trustzone, *Arm TrustZone Technology*, <https://developer.arm.com/ip-products/security-ip/trustzone> (accessed Dec. 5, 2019).
- [256] Android trusty, *Android trusty TEE*, <https://source.android.com/security/trusty> (accessed Dec. 5, 2019).

- [257] V. Costan, I. Lebedev, and S. Devadas, “Sanctum: Minimal hardware extensions for strong software isolation,” in *25th USENIX Security Symposium (USENIX Security 16)*, pp. 857–874, 2016.
- [258] F. Tramèr and D. Boneh, “Slalom: Fast, verifiable and private execution of neural networks in trusted hardware,” in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=rJVorjCcKQ>.
- [259] J. Van Bulck, M. Minkin, O. Weisse, D. Genkin, B. Kasikci, F. Piessens, M. Silberstein, T. F. Wenisch, Y. Yarom, and R. Strackx, “Foreshadow: Extracting the keys to the intel SGX kingdom with transient out-of-order execution,” in *27th USENIX Security Symposium (USENIX Security 18)*, pp. 991–1008, 2018.
- [260] G. Ács and C. Castelluccia, “I have a DREAM!: Differentially privatE smArt Metering,” in *Proceedings of the 13th International Conference on Information Hiding, IH’11*, pp. 118–132, Berlin, Heidelberg: Springer-Verlag, 2011. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2042445.2042457>.
- [261] J. So, B. Guler, and A. S. Avestimehr, “Turbo-aggregate: Breaking the quadratic aggregation barrier in secure federated learning,” *arXiv preprint arXiv:2002.04156*, 2020.
- [262] S. Goryczka and L. Xiong, “A comprehensive comparison of multiparty secure additions with differential privacy,” *IEEE Trans. Dependable Sec. Comput.*, vol. 14, no. 5, pp. 463–477, 2017.
- [263] T.-H. H. Chan, E. Shi, and D. Song, “Privacy-preserving stream aggregation with fault tolerance,” in *International Conference on Financial Cryptography and Data Security*, Springer, pp. 200–214, 2012.
- [264] S. Halevi, Y. Lindell, and B. Pinkas, “Secure computation on the web: Computing without simultaneous interaction,” in *Annual Cryptology Conference*, Springer, pp. 132–150, 2011.
- [265] E. Shi, H. T. H. Chan, E. Rieffel, R. Chow, and D. Song, “Privacy-preserving aggregation of time-series data,” in *Annual Network & Distributed System Security Symposium (NDSS)*, 2011.

- [266] M. Burkhart, M. Strasser, D. Many, and X. Dimitropoulos, “SEPIA: Privacy-preserving aggregation of multi-domain network events and statistics,” *Network*, vol. 1, no. 101101, 2010.
- [267] H. Corrigan-Gibbs and D. Boneh, “Prio: Private, robust, and scalable computation of aggregate statistics,” in *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, pp. 259–282, 2017.
- [268] D. Lie and P. Maniatis, “Glimmers: Resolving the privacy/trust quagmire,” in *Proceedings of the 16th Workshop on Hot Topics in Operating Systems*, ACM, pp. 94–99, 2017.
- [269] D. Chaum, “Untraceable electronic mail, return addresses, and digital pseudonyms,” *Communications of the ACM*, vol. 24, no. 2, 1981.
- [270] A. Kwon, D. Lazar, S. Devadas, and B. Ford, “Riffle,” *Proceedings on Privacy Enhancing Technologies*, vol. 2016, no. 2, pp. 115–134, 2016.
- [271] R. Dingledine, N. Mathewson, and P. Syverson, *Tor: The second-generation onion router*, Tech. Rep., Naval Research Lab, Washington, DC, 2004.
- [272] E. Kushilevitz and R. Ostrovsky, “Replication is not needed: Single database, computationally-private information retrieval,” in *Proc. of the 38th Annu. IEEE Symp. on Foundations of Computer Science*, pp. 364–373, 1997.
- [273] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan, “Private information retrieval,” *J. ACM*, vol. 45, no. 6, pp. 965–981, Nov. 1998.
- [274] R. Sion and B. Carbunar, “On the computational practicality of private information retrieval,” in *Proceedings of the Network and Distributed Systems Security Symposium*, Internet Society, pp. 2006–06, 2007.
- [275] C. Aguilar-Melchor and P. Gaborit, “A lattice-based computationally-efficient private information retrieval protocol,” *Cryptol. ePrint Arch., Report*, vol. 446, 2007.

- [276] C. Aguilar-Melchor, J. Barrier, L. Fousse, and M.-O. Killijian, “XPIR: Private information retrieval for everyone,” *Proceedings on Privacy Enhancing Technologies*, vol. 2016, no. 2, pp. 155–174, 2016.
- [277] S. Angel, H. Chen, K. Laine, and S. T. V. Setty, “PIR with compressed queries and amortized query processing,” in *IEEE Symposium on Security and Privacy*, pp. 962–979, IEEE Computer Society, 2018.
- [278] C. Gentry and S. Halevi, “Compressible FHE with applications to PIR,” in *TCC (2)*, Lecture Notes in Computer Science, vol. 11892, pp. 438–464, Springer, 2019.
- [279] F. Olumofin and I. Goldberg, “Revisiting the computational practicality of private information retrieval,” in *International Conference on Financial Cryptography and Data Security*, Springer, pp. 158–172, 2011.
- [280] A. Ali, T. Lepoint, S. Patel, M. Raykova, P. Schoppmann, K. Seth, and K. Yeo, “Communication-computation trade-offs in PIR,” *IACR Cryptol. ePrint Arch.*, vol. 2019, p. 1483, 2019.
- [281] S. Kadhe, B. Garcia, A. Heidarzadeh, S. E. Rouayheb, and A. Sprintson, “Private information retrieval with side information: The single server case,” in *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 1099–1106, Oct. 2017. DOI: [10.1109/ALLERTON.2017.8262860](https://doi.org/10.1109/ALLERTON.2017.8262860).
- [282] S. Patel, G. Persiano, and K. Yeo, “Private stateful information retrieval,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS '18*, pp. 1002–1019, New York, NY, USA: ACM, 2018. DOI: [10.1145/3243734.3243821](https://doi.org/10.1145/3243734.3243821).
- [283] H. Corrigan-Gibbs and D. Kogan, “Private information retrieval with sublinear online time,” *IACR Cryptology ePrint Archive*, vol. 2019, p. 1075, 2019.
- [284] D. Woodruff and S. Yekhanin, “A geometric approach to information-theoretic private information retrieval,” in *20th Annual IEEE Conference on Computational Complexity (CCC'05)*, pp. 275–284, Jun. 2005. DOI: [10.1109/CCC.2005.2](https://doi.org/10.1109/CCC.2005.2).

- [285] R. G. L. D'Oliveira and S. E. Rouayheb, "Lifting private information retrieval from two to any number of messages," in *2018 IEEE International Symposium on Information Theory (ISIT)*, pp. 1744–1748, Jun. 2018. DOI: [10.1109/ISIT.2018.8437805](https://doi.org/10.1109/ISIT.2018.8437805).
- [286] R. Bitar and S. E. Rouayheb, "Staircase-PIR: Universally robust private information retrieval," in *2018 IEEE Information Theory Workshop (ITW)*, pp. 1–5, Nov. 2018. DOI: [10.1109/ITW.2018.8613532](https://doi.org/10.1109/ITW.2018.8613532).
- [287] T. Lepoint, S. Patel, M. Raykova, K. Seth, and N. Trieu, "Private join and compute from PIR with default," *IACR Cryptol. ePrint Arch.*, vol. 2020, p. 1011, 2020.
- [288] C. Naim, F. Ye, and S. E. Rouayheb, "ON-OFF privacy with correlated requests," in *2019 IEEE International Symposium on Information Theory (ISIT)*, Jul. 2019.
- [289] F. Ye, C. Naim, and S. El Rouayheb, "Preserving ON-OFF privacy for past and future requests," in *2019 IEEE Information Theory Workshop (ITW)*, Aug. 2019.
- [290] C. Dwork, F. McSherry, K. Nissim, and A. D. Smith, "Calibrating noise to sensitivity in private data analysis," in *IACR Theory of Cryptography Conference (TCC), New York*, Lecture Notes in Computer Science, vol. 3876, pp. 265–284, Springer-Verlag, 2006. DOI: [10.1007/11681878_14](https://doi.org/10.1007/11681878_14).
- [291] C. Dwork, "Differential privacy: A survey of results," in *International Conference on Theory and Applications of Models of Computation*, Springer, pp. 1–19, 2008.
- [292] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations and Trends in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.
- [293] S. L. Warner, "Randomized response: A survey technique for eliminating evasive answer bias," *J. Am. Stat. Assoc.*, vol. 60, no. 309, pp. 63–69, 1965.
- [294] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. D. Smith, "What can we learn privately?" *SIAM J. Comput.*, vol. 40, no. 3, pp. 793–826, 2011.

- [295] B. Ding, J. Kulkarni, and S. Yekhanin, “Collecting telemetry data privately,” in *Advances in Neural Information Processing Systems*, vol. 30, Dec. 2017. [Online]. Available: <https://www.microsoft.com/en-us/research/publication/collecting-telemetry-data-privately/>.
- [296] V. Pihur, A. Korolova, F. Liu, S. Sankuratripati, M. Yung, D. Huang, and R. Zeng, “Differentially-private ‘draw and discard’ machine learning,” *CoRR*, vol. abs/1807.04369, 2018. [Online]. Available: <http://arxiv.org/abs/1807.04369>.
- [297] J. Ullman, *Tight lower bounds for locally differentially private selection*, Tech. Rep., abs/1802.02638, 2018. [Online]. Available: <http://arxiv.org/abs/1802.02638>.
- [298] A. Cheu, A. Smith, J. Ullman, D. Zeber, and M. Zhilyaev, “Distributed differential privacy via shuffling,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Springer, pp. 375–403, 2019.
- [299] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, “Our data, ourselves: Privacy via distributed noise generation,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Springer, pp. 486–503, 2006.
- [300] C. Sabater, A. Bellet, and J. Ramon, “Distributed differentially private averaging with improved utility and robustness to malicious parties,” *arXiv preprint arXiv:2006.07218*, 2020.
- [301] A. Ghosh, T. Roughgarden, and M. Sundararajan, “Universally utility-maximizing privacy mechanisms,” in *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC ’09, pp. 351–360, New York, NY, USA: ACM, 2009. [Online]. Available: <http://doi.acm.org/10.1145/1536414.1536464>.
- [302] V. Rastogi and S. Nath, “Differentially private aggregation of distributed time-series with transformation and encryption,” in *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*, SIGMOD ’10, pp. 735–746, New York, NY, USA: ACM, 2010. [Online]. Available: <http://doi.acm.org/10.1145/1807167.1807247>.

- [303] Ú. Erlingsson, V. Feldman, I. Mironov, A. Raghunathan, K. Talwar, and A. Thakurta, “Amplification by shuffling: From local to central differential privacy via anonymity,” in *SODA*, pp. 2468–2479, 2019.
- [304] B. Balle, J. Bell, A. Gascón, and K. Nissim, “The privacy blanket of the shuffle model,” in *Advances in Cryptology – CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2019, Proceedings, Part II*, pp. 638–667, 2019. DOI: [10.1007/978-3-030-26951-7_22](https://doi.org/10.1007/978-3-030-26951-7_22).
- [305] L. Chen, B. Ghazi, R. Kumar, and P. Manurangsi, “On distributed differential privacy and counting distinct elements,” in *Innovations in Theoretical Computer Science (ITCS)*, 2021.
- [306] B. Ghazi, R. Kumar, P. Manurangsi, and R. Pagh, “Private counting from anonymous messages: Near-optimal accuracy with vanishing communication overhead,” in *ICML*, 2020.
- [307] B. Ghazi, R. Pagh, and A. Velingker, “Scalable and differentially private distributed aggregation in the shuffled model,” *arXiv preprint arXiv:1906.08320*, 2019.
- [308] B. Ghazi, P. Manurangsi, R. Pagh, and A. Velingker, “Private aggregation from fewer anonymous messages,” in *EUROCRYPT*, pp. 798–827, 2020.
- [309] B. Ghazi, N. Golowich, R. Kumar, R. Pagh, and A. Velingker, “On the power of multiple anonymous messages,” *arXiv:1908.11358*, 2019.
- [310] B. Ghazi, N. Golowich, R. Kumar, P. Manurangsi, R. Pagh, and A. Velingker, “Pure differentially private summation from anonymous messages,” in *ITC*, pp. 15:1–15:23, 2020.
- [311] B. Avent, A. Korolova, D. Zeber, T. Hovden, and B. Livshits, “BLENDER: Enabling local search with a hybrid differential privacy model,” in *26th USENIX Security Symposium (USENIX Security 17)*, pp. 747–764, Vancouver, BC: USENIX Association, Aug. 2017. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/avent>.

- [312] A. Beimel, A. Korolova, K. Nissim, O. Sheffet, and U. Stemmer, “The power of synergy in differential privacy: Combining a small curator with local randomizers,” in *Conference on Information-Theoretic Cryptography (ITC)*, 2020. [Online]. Available: <https://arxiv.org/abs/1912.08951>.
- [313] L. Babai, L. Fortnow, L. A. Levin, and M. Szegedy, “Checking computations in polylogarithmic time,” in *STOC*, pp. 21–31, ACM, 1991.
- [314] S. Micali, “Computationally sound proofs,” *SIAM J. Comput.*, vol. 30, no. 4, pp. 1253–1298, 2000.
- [315] S. Goldwasser, Y. T. Kalai, and G. N. Rothblum, “Delegating computation: Interactive proofs for muggles,” in *STOC*, pp. 113–122, ACM, 2008.
- [316] R. Gennaro, C. Gentry, and B. Parno, “Non-interactive verifiable computing: Outsourcing computation to untrusted workers,” in *CRYPTO*, Lecture Notes in Computer Science, vol. 6223, pp. 465–482, Springer, 2010.
- [317] S. Goldwasser, S. Micali, and C. Rackoff, “The knowledge complexity of interactive proof systems,” *SIAM J. Comput.*, vol. 18, no. 1, pp. 186–208, 1989.
- [318] B. Parno, J. Howell, C. Gentry, and M. Raykova, “Pinocchio: Nearly practical verifiable computation,” *Commun. ACM*, vol. 59, no. 2, pp. 103–112, 2016.
- [319] C. P. Schnorr, “Efficient identification and signatures for smart cards,” in *Proceedings of the Workshop on the Theory and Application of Cryptographic Techniques on Advances in Cryptology*, EUROCRYPT ’89, 1990.
- [320] Damgård, *On σ protocols*, <http://www.cs.au.dk/~ivan/Sigma.pdf>, 2010.
- [321] N. Bitansky, R. Canetti, A. Chiesa, and E. Tromer, “From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again,” in *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS ’12, 2012.

- [322] R. Gennaro, C. Gentry, B. Parno, and M. Raykova, “Quadratic span programs and succinct NIZKs without PCPs,” in *EUROCRYPT*, Lecture Notes in Computer Science, vol. 7881, pp. 626–645, Springer, 2013.
- [323] C. Costello, C. Fournet, J. Howell, M. Kohlweiss, B. Kreuter, M. Naehrig, B. Parno, and S. Zahur, “Geppetto: Versatile verifiable computation,” in *IEEE Symposium on Security and Privacy*, pp. 253–270, IEEE Computer Society, 2015.
- [324] *libsnark: A c++ library for zkSNARK proofs*, <https://github.com/scipr-lab/libsnark>, Dec. 2019.
- [325] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza, “Zerocash: Decentralized anonymous payments from bitcoin,” in *IEEE Symposium on Security and Privacy*, pp. 459–474, IEEE Computer Society, 2014.
- [326] T. Xie, J. Zhang, Y. Zhang, C. Papamanthou, and D. Song, “Libra: Succinct zero-knowledge proofs with optimal prover computation,” in *CRYPTO (3)*, Lecture Notes in Computer Science, vol. 11694, pp. 733–764, Springer, 2019.
- [327] S. Ames, C. Hazay, Y. Ishai, and M. Venkatasubramanian, “Ligero: Lightweight sublinear arguments without a trusted setup,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS ’17*, 2017.
- [328] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell, “Bulletproofs: Short proofs for confidential transactions and more,” in *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21–23 May 2018, San Francisco, California, USA*, 2018.
- [329] R. S. Wahby, I. Tzialla, A. Shelat, J. Thaler, and M. Walfish, “Doubly-efficient zksnarks without trusted setup,” in *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21–23 May 2018, San Francisco, California, USA*, 2018.
- [330] E. Ben-Sasson, I. Bentov, Y. Horesh, and M. Riabzev, “Scalable zero knowledge with no trusted setup,” in *CRYPTO (3)*, Lecture Notes in Computer Science, vol. 11694, pp. 701–732, Springer, 2019.

- [331] D. Boneh, E. Boyle, H. Corrigan-Gibbs, N. Gilboa, and Y. Ishai, “Zero-knowledge proofs on secret-shared data via fully linear PCPs,” in *CRYPTO (3)*, Lecture Notes in Computer Science, vol. 11694, pp. 67–97, Springer, 2019.
- [332] F. Tramèr, F. Zhang, H. Lin, J. Hubaux, A. Juels, and E. Shi, “Sealed-glass proofs: Using transparent enclaves to prove and sell knowledge,” in *2017 IEEE European Symposium on Security and Privacy, EuroS&P 2017, Paris, France, April 26–28, 2017*, pp. 19–34, 2017.
- [333] A. Seshadri, M. Luk, A. Perrig, L. van Doorn, and P. K. Khosla, “Pioneer: Verifying code integrity and enforcing untampered code execution on legacy systems,” in *Malware Detection, Advances in Information Security*, vol. 27, Springer, 2007, pp. 253–289.
- [334] K. Eldefrawy, G. Tsudik, A. Francillon, and D. Perito, “SMART: Secure and minimal architecture for (establishing dynamic) root of trust,” in *NDSS*, The Internet Society, 2012.
- [335] P. Koeberl, S. Schulz, A.-R. Sadeghi, and V. Varadharajan, “TrustLite: A security architecture for tiny embedded devices,” in *EuroSys*, 10:1–10:14, ACM, 2014.
- [336] A. Francillon, Q. Nguyen, K. B. Rasmussen, and G. Tsudik, “A minimalist approach to remote attestation,” in *DATE*, pp. 1–6, European Design and Automation Association, 2014.
- [337] N. Carlini, C. Liu, J. Kos, Ú. Erlingsson, and D. Song, “The secret sharer: Measuring unintended neural network memorization and extracting secrets,” *arXiv preprint arXiv:1802.08232*, 2018.
- [338] N. Carlini, F. Tramèr, E. Wallace, M. Jagielski, A. Herbert-Voss, K. Lee, A. Roberts, T. Brown, D. Song, Ú. Erlingsson, A. Oprea, and C. Raffel, “Extracting training data from large language models,” *arXiv preprint arXiv:2012.07805*, 2020.
- [339] M. Fredrikson, S. Jha, and T. Ristenpart, “Model inversion attacks that exploit confidence information and basic countermeasures,” in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, ACM, pp. 1322–1333, 2015.

- [340] L. Melis, C. Song, E. De Cristofaro, and V. Shmatikov, “Exploiting unintended feature leakage in collaborative learning,” *arXiv preprint arXiv:1805.04049*, 2018.
- [341] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, “Membership inference attacks against machine learning models,” in *2017 IEEE Symposium on Security and Privacy (SP)*, IEEE, pp. 3–18, 2017.
- [342] S. Yeom, I. Giacomelli, M. Fredrikson, and S. Jha, “Privacy risk in machine learning: Analyzing the connection to overfitting,” in *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*, IEEE, pp. 268–282, 2018.
- [343] M. Diaz, P. Kairouz, J. Liao, and L. Sankar, “Theoretical guarantees for model auditing with finite adversaries,” *arXiv preprint arXiv:1911.03405*, 2019.
- [344] A. Bhowmick, J. Duchi, J. Freudiger, G. Kapoor, and R. Rogers, “Protection against reconstruction and its applications in private federated learning,” *arXiv preprint arXiv:1812.00984*, 2018.
- [345] H. B. McMahan, G. Andrew, Ú. Erlingsson, S. Chien, I. Mironov, N. Papernot, and P. Kairouz, “A general approach to adding differential privacy to iterative training procedures,” *arXiv preprint arXiv:1812.06210*, 2018.
- [346] C. Dwork, G. N. Rothblum, and S. Vadhan, “Boosting and differential privacy,” in *Proceedings of the IEEE 51st Annual Symposium on Foundations of Computer Science, FOCS '10*, pp. 51–60, 2010.
- [347] P. Kairouz, S. Oh, and P. Viswanath, “The composition theorem for differential privacy,” *IEEE Trans. Inform. Theor.*, vol. 63, no. 6, pp. 4037–4049, 2017.
- [348] I. Mironov, “Rényi differential privacy,” in *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, IEEE, pp. 263–275, 2017.
- [349] I. Mironov, K. Talwar, and L. Zhang, “Rényi differential privacy of the sampled Gaussian mechanism,” *arXiv preprint arXiv:1908.10530*, 2019.
- [350] Y.-X. Wang, B. Balle, and S. Kasiviswanathan, “Subsampled Rényi differential privacy and analytical moments accountant,” *arXiv preprint arXiv:1808.00087*, 2018.

- [351] S. Ramaswamy, O. Thakkar, R. Mathews, G. Andrew, H. B. McMahan, and F. Beaufays, “Training production language models without memorizing user data,” *arXiv preprint arXiv:2009.10031*, 2020.
- [352] N. Papernot, A. Thakurta, S. Song, S. Chien, and Ú. Erlingsson, “Tempered sigmoid activations for deep learning with differential privacy,” *arXiv preprint arXiv:2007.14191*, 2020.
- [353] F. Tramèr and D. Boneh, “Differentially private learning needs better features (or much more data),” *arXiv preprint arXiv:2011.11660*, 2020.
- [354] K. Amin, A. Kulesza, A. Munoz, and S. Vassilvtiskii, “Bounding user contributions: A bias-variance trade-off in differential privacy,” in *International Conference on Machine Learning*, pp. 263–271, 2019.
- [355] V. Pichapati, A. T. Suresh, F. X. Yu, S. J. Reddi, and S. Kumar, “AdaClip: Adaptive clipping for private SGD,” *arXiv preprint arXiv:1908.07643*, 2019.
- [356] Y. Liu, A. T. Suresh, F. X. X. Yu, S. Kumar, and M. Riley, “Learning discrete distributions: User vs item-level privacy,” in *Advances in Neural Information Processing Systems*, vol. 33, pp. 20 965–20 976, 2020.
- [357] V. Feldman, I. Mironov, K. Talwar, and A. Thakurta, “Privacy amplification by iteration,” in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, pp. 521–532, 2018.
- [358] B. Balle, P. Kairouz, H. B. McMahan, O. Thakkar, and A. Thakurta, “Privacy amplification via random check-ins,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, Eds., vol. 33, pp. 4623–4634, Curran Associates, Inc., 2020.
- [359] P. Kairouz, B. McMahan, S. Song, O. Thakkar, A. Thakurta, and Z. Xu, *Practical and private (deep) learning without sampling or shuffling*, 2021.

- [360] J. K. Salmon, M. A. Moraes, R. O. Dror, and D. E. Shaw, “Parallel random numbers: As easy as 1, 2, 3,” in *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, ACM, p. 16, 2011.
- [361] I. Mironov, O. Pandey, O. Reingold, and S. Vadhan, “Computational differential privacy,” in *Advances in Cryptology—CRYPTO*, pp. 126–142, 2009.
- [362] A. Haeberlen, B. C. Pierce, and A. Narayan, “Differential privacy under fire,” in *USENIX Security Symposium*, 2011.
- [363] I. Mironov, “On significance of the least significant bits for differential privacy,” in *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, ACM, pp. 650–661, 2012.
- [364] Z. Ding, Y. Wang, G. Wang, D. Zhang, and D. Kifer, “Detecting violations of differential privacy,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS ’18*, pp. 475–489, New York, NY, USA: ACM, 2018. [Online]. Available: <http://doi.acm.org/10.1145/3243734.3243818>.
- [365] M. Jagielski, J. Ullman, and A. Oprea, “Auditing differentially private machine learning: How private is private SGD?” in *Advances in Neural Information Processing Systems*, vol. 33, pp. 22 205–22 216, 2020.
- [366] X. Liu and S. Oh, “Minimax rates of estimating approximate differential privacy,” *arXiv preprint arXiv:1905.10335*, 2019.
- [367] M. S. Riazi, K. Laine, B. Pelton, and W. Dai, “HEAX: High-performance architecture for computation on homomorphically encrypted data in the cloud,” *arXiv preprint arXiv:1909.09731*, 2019.
- [368] R. Cummings, S. Krehbiel, Y. Mei, R. Tuo, and W. Zhang, “Differentially private change-point detection,” in *Advances in Neural Information Processing Systems, NeurIPS ’18*, vol. 31, pp. 10 825–10 834, 2018.
- [369] R. Cummings, S. Krehbiel, K. Lai, and U. Tantitongpipat, “Differential privacy for growing databases,” in *Advances in Neural Information Processing Systems, NeurIPS ’18*, vol. 31, pp. 8864–8873, 2018.

- [370] C. L. Canonne, G. Kamath, A. McMillan, A. Smith, and J. Ullman, “The structure of optimal private tests for simple hypotheses,” *arXiv preprint arXiv:1811.11148*, 2019.
- [371] N. C. Abay, Y. Zhou, M. Kantarcioglu, B. Thuraisingham, and L. Sweeney, “Privacy preserving synthetic data release using deep learning,” in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, Springer, pp. 510–526, 2018.
- [372] F. Tramèr, F. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, “Stealing machine learning models via prediction APIs,” in *25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10–12, 2016.*, pp. 601–618, 2016. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/tramer>.
- [373] J. R. Douceur, “The sybil attack,” in *Revised Papers from the First International Workshop on Peer-to-Peer Systems, IPTPS ’01*, pp. 251–260, London, UK: Springer-Verlag, 2002. [Online]. Available: <http://dl.acm.org/citation.cfm?id=646334.687813>.
- [374] R. Bassily, N. Kobbi, U. Stemmer, and A. G. Thakurta, “Practical locally private heavy hitters,” *J. Mach. Learn. Res.*, vol. 21, no. 16, pp. 1–42, 2020.
- [375] G. Cormode, T. Kulkarni, and D. Srivastava, “Marginal release under local differential privacy,” in *Proceedings of the 2018 International Conference on Management of Data*, ACM, pp. 131–146, 2018.
- [376] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, “Local privacy and statistical minimax rates,” in *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, IEEE, pp. 429–438, 2013.
- [377] P. Kairouz, S. Oh, and P. Viswanath, “Extremal mechanisms for local differential privacy,” in *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds., vol. 27, pp. 2879–2887, Curran Associates, Inc., 2014.

- [378] P. Kairouz, K. A. Bonawitz, and D. Ramage, “Discrete distribution estimation under local privacy,” in *International Conference on Machine Learning*, pp. 2436–2444, 2016.
- [379] M. Ye and A. Barg, “Optimal schemes for discrete distribution estimation under locally differential privacy,” *IEEE Trans. Inform. Theor.*, 2018.
- [380] J. Tang, A. Korolova, X. Bai, X. Wang, and X. Wang, “Privacy loss in Apple’s implementation of differential privacy on MacOS 10.12,” *CoRR*, vol. abs/1709.02753, 2017. [Online]. Available: <http://arxiv.org/abs/1709.02753>.
- [381] B. Avent, Y. Dubey, and A. Korolova, “The power of the hybrid model for mean estimation,” *Proceedings on Privacy Enhancing Technologies (PETS)*, vol. 2020, no. 4, pp. 48–68, 1 Oct. 2020. DOI: <https://doi.org/10.2478/popets-2020-0062>.
- [382] E. Kushilevitz and N. Nisan, *Communication Complexity*. New York, NY, USA: Cambridge University Press, 1997.
- [383] W.-N. Chen, P. Kairouz, and A. Ozgur, “Breaking the communication-privacy-accuracy trilemma,” in *Advances in Neural Information Processing Systems*, vol. 33, pp. 3312–3324, 2020.
- [384] B. Balle, J. Bell, A. Gascón, and K. Nissim, “Private summation in the multi-message shuffle model,” in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, ACM, pp. 657–676, 2020.
- [385] R. Bassily and A. Smith, “Local, private, efficient protocols for succinct histograms,” in *STOC*, pp. 127–135, 2015.
- [386] F. McSherry and K. Talwar, “Mechanism design via differential privacy,” in *FOCS*, pp. 94–103, 2007.
- [387] T. Steinke and J. Ullman, “Tight lower bounds for differentially private selection,” in *FOCS*, pp. 552–563, 2017.
- [388] A. M. Girgis, D. Data, S. Diggavi, P. Kairouz, and A. T. Suresh, “Shuffled model of federated learning: Privacy, communication and accuracy trade-offs,” *arXiv preprint arXiv:2008.07180*, 2020.
- [389] P. Kairouz, Z. Liu, and T. Steinke, “The distributed discrete gaussian mechanism for federated learning with secure aggregation,” *arXiv preprint arXiv:2102.06387*, 2021.

- [390] W.-T. Chang and R. Tandon, “On the upload versus download cost for secure and private matrix multiplication,” *ArXiv, abs/1906.10684*, 2019.
- [391] Z. Jia and S. A. Jafar, “On the capacity of secure distributed matrix multiplication,” *CoRR*, vol. abs/1908.06957, 2019.
- [392] C. Niu, F. Wu, S. Tang, L. Hua, R. Jia, C. Lv, Z. Wu, and G. Chen, “Secure federated submodel learning,” *arXiv preprint arXiv:1911.02254*, 2019.
- [393] M. J. Kearns, A. Roth, Z. S. Wu, and G. Yaroslavtsev, “Privacy for the protected (only),” *CoRR*, vol. abs/1506.00242, 2015. [Online]. Available: <http://arxiv.org/abs/1506.00242>.
- [394] D. Kifer and A. Machanavajjhala, “Pufferfish: A framework for mathematical privacy definitions,” *ACM Trans. Database Sys.*, vol. 39, no. 1, 3:1–3:36, 2014.
- [395] J. M. Abowd and I. M. Schmutte, “An economic analysis of privacy protection and statistical accuracy as social choices,” *Am. Econ. Rev.*, vol. 109, no. 1, pp. 171–202, 2019.
- [396] R. Cummings, I. Dekel, O. Heffetz, and K. Ligett, “Bringing differential privacy into the experimental economics lab: Theory and an application to a public-good game,” Working paper, 2019.
- [397] B. Biggio, B. Nelson, and P. Laskov, “Poisoning attacks against support vector machines,” in *Proceedings of the 29th International Conference on International Conference on Machine Learning, ICML’12*, pp. 1467–1474, Edinburgh, UK: Omnipress, 2012. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3042573.3042761>.
- [398] Y. Liu, S. Ma, Y. Aafer, W.-C. Lee, J. Zhai, W. Wang, and X. Zhang, “Trojaning attack on neural networks,” in *25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18–21, 2018*, 2018.
- [399] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, “How to backdoor federated learning,” *arXiv preprint arXiv:1807.00459*, 2018.

- [400] A. N. Bhagoji, S. Chakraborty, P. Mittal, and S. Calo, “Analyzing federated learning through an adversarial lens,” in *Proceedings of the 36th International Conference on Machine Learning*, pp. 634–643, 2019.
- [401] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings*, 2015. [Online]. Available: <http://arxiv.org/abs/1412.6572>.
- [402] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, “Intriguing properties of neural networks,” *ICLR*, 2013.
- [403] X. Chen, C. Liu, B. Li, K. Lu, and D. Song, “Targeted backdoor attacks on deep learning systems using data poisoning,” *arXiv preprint arXiv:1712.05526*, 2017.
- [404] D. Alistarh, Z. Allen-Zhu, and J. Li, “Byzantine stochastic gradient descent,” in *NIPS*, 2018.
- [405] P. Blanchard, E. M. El Mahdi, R. Guerraoui, and J. Stainer, “Machine learning with adversaries: Byzantine tolerant gradient descent,” in *Advances in Neural Information Processing Systems*, pp. 118–128, 2017.
- [406] Y. Chen, L. Su, and J. Xu, “Distributed statistical machine learning in adversarial settings: Byzantine gradient descent,” *POMACS*, vol. 1, pp. 44:1–44:25, 2017.
- [407] L. Chen, H. Wang, Z. B. Charles, and D. S. Papailiopoulos, “DRACO: Byzantine-resilient distributed training via redundant gradients,” in *Proceedings of the 35th International Conference on Machine Learning, ICML*, 2018.
- [408] E. M. El Mhamdi, R. Guerraoui, and S. Rouault, “The hidden vulnerability of distributed learning in Byzantium,” in *ICML*, 2018.
- [409] G. F. Cretu, A. Stavrou, M. E. Locasto, S. J. Stolfo, and A. D. Keromytis, “Casting out demons: Sanitizing training data for anomaly sensors,” in *2008 IEEE Symposium on Security and Privacy (sp 2008)*, IEEE, pp. 81–95, 2008.

- [410] J. Steinhardt, P. W. W. Koh, and P. S. Liang, “Certified defenses for data poisoning attacks,” in *Advances in Neural Information Processing Systems*, pp. 3517–3529, 2017.
- [411] I. Diakonikolas, G. Kamath, D. Kane, J. Li, J. Steinhardt, and A. Stewart, “Sever: A robust meta-algorithm for stochastic optimization,” in *Proceedings of the 36th International Conference on Machine Learning*, Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97, pp. 1596–1606, Long Beach, California, USA: PMLR 9–15 Jun. 2019, Sep. 2019. [Online]. Available: <http://proceedings.mlr.press/v97/diakonikolas19a.html>.
- [412] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrndić, P. Laskov, G. Giacinto, and F. Roli, “Evasion attacks against machine learning at test time,” in *ECML-PKDD*, Springer, pp. 387–402, 2013.
- [413] N. Carlini and D. Wagner, “Towards evaluating the robustness of neural networks,” in *2017 IEEE Symposium on Security and Privacy (SP)*, IEEE, pp. 39–57, 2017.
- [414] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” *ICLR*, 2015.
- [415] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards deep learning models resistant to adversarial attacks,” *ICLR*, 2017.
- [416] L. Lamport, R. Shostak, and M. Pease, “The Byzantine generals problem,” *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 382–401, 1982.
- [417] D. Yin, Y. Chen, K. Ramchandran, and P. Bartlett, “Byzantine-robust distributed learning: Towards optimal statistical rates,” in *ICML*, 2019.
- [418] L. Su and N. H. Vaidya, “Fault-Tolerant Multi-Agent Optimization: Optimal Iterative Distributed Algorithms,” in *PODC*, 2016.
- [419] J. So, B. Guler, and A. S. Avestimehr, “Byzantine-resilient secure federated learning,” in *IEEE Journal on Selected Areas in Communication, Series on Machine Learning for Communications and Networks*, 2020.

- [420] K. Pillutla, S. M. Kakade, and Z. Harchaoui, “Robust aggregation for federated learning,” *arXiv preprint arXiv:1912.13445*, 2019.
- [421] C. Xie, S. Koyejo, and I. Gupta, “Practical distributed learning: Secure machine learning with communication-efficient local updates,” in *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD)*, 2019.
- [422] M. Fang, X. Cao, J. Jia, and N. Z. Gong, “Local model poisoning attacks to Byzantine-robust federated learning,” *arXiv preprint arXiv:1911.11815*, 2019.
- [423] M. Baruch, G. Baruch, and Y. Goldberg, “A little is enough: Circumventing defenses for distributed learning,” *arXiv preprint arXiv:1902.06156*, 2019.
- [424] S. Rajput, H. Wang, Z. Charles, and D. Papailiopoulos, “DETOX: A redundancy-based framework for faster and more robust gradient aggregation,” *arXiv preprint arXiv:1907.12205*, 2019.
- [425] D. Data, L. Song, and S. Diggavi, “Data encoding for byzantine-resilient distributed optimization,” *IEEE Trans. Inform. Theor.*, pp. 2719–2723, 2020.
- [426] C.-L. Chen, L. Golubchik, and M. Paolieri, “Backdoor attacks on federated meta-learning,” *arXiv preprint arXiv:2006.07026*, 2020.
- [427] E. Chou, F. Tramèr, and G. Pellegrino, “SentiNet: Detecting physical attacks against deep learning systems,” *arXiv preprint arXiv:1812.00292*, 2018.
- [428] K. Liu, B. Dolan-Gavitt, and S. Garg, “Fine-pruning: Defending against backdooring attacks on deep neural networks,” in *International Symposium on Research in Attacks, Intrusions, and Defenses*, Springer, pp. 273–294, 2018.
- [429] Y. Shen and S. Sanghavi, “Learning with bad training data via iterative trimmed loss minimization,” in *Proceedings of the 36th International Conference on Machine Learning*, Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97, pp. 5739–5748, Long Beach, California, USA: PMLR. 9–15 Jun. 2019, 2019. [Online]. Available: <http://proceedings.mlr.press/v97/shen19e.html>.

- [430] B. Tran, J. Li, and A. Madry, “Spectral signatures in backdoor attacks,” in *Advances in Neural Information Processing Systems*, pp. 8000–8010, 2018.
- [431] B. Wang, Y. Yao, S. Shan, H. Li, B. Viswanath, H. Zheng, and B. Y. Zhao, “Neural cleanse: Identifying and mitigating backdoor attacks in neural networks,” in *2019 IEEE Symposium on Security and Privacy*, IEEE, 2019.
- [432] C. Fung, C. J.-M. Yoon, and I. Beschastnikh, “Mitigating sybils in federated learning poisoning,” *arXiv preprint arXiv:1808.04866*, 2018.
- [433] P. W. Koh and P. Liang, “Understanding black-box predictions via influence functions,” in *Proceedings of the 34th International Conference on Machine Learning – Volume 70*, JMLR.org, pp. 1885–1894, 2017.
- [434] C. Xie, S. Koyejo, and I. Gupta, “Zeno: Distributed stochastic gradient descent with suspicion-based fault-tolerance,” in *International Conference on Machine Learning*, pp. 6893–6901, 2019.
- [435] C. Xie, “Zeno++: Robust asynchronous SGD with arbitrary number of Byzantine workers,” *arXiv preprint arXiv:1903.07020*, 2019.
- [436] T. Gu, B. Dolan-Gavitt, and S. Garg, “BadNets: Identifying vulnerabilities in the machine learning model supply chain,” *arXiv preprint arXiv:1708.06733*, 2017.
- [437] P. W. Koh, J. Steinhardt, and P. Liang, “Stronger data poisoning attacks break data sanitization defenses,” *arXiv preprint arXiv:1811.00741*, 2018.
- [438] X. Zhu, “Machine teaching: An inverse problem to machine learning and an approach toward optimal education,” in *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [439] S. Mei and X. Zhu, “Using machine teaching to identify optimal training-set attacks on machine learners,” in *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [440] L. Engstrom, B. Tran, D. Tsipras, L. Schmidt, and A. Madry, “A rotation and a translation suffice: Fooling CNNs with simple transformations,” *arXiv preprint arXiv:1712.02779*, 2017.

- [441] D. Kang, Y. Sun, D. Hendrycks, T. Brown, and J. Steinhardt, “Testing robustness against unforeseen adversaries,” *arXiv preprint arXiv:1908.08016*, 2019.
- [442] E. Wong, F. R. Schmidt, and J. Z. Kolter, “Wasserstein adversarial examples via projected sinkhorn iterations,” *ICML*, 2019.
- [443] A. Kurakin, I. Goodfellow, and S. Bengio, “Adversarial machine learning at scale,” *arXiv preprint arXiv:1611.01236*, 2016.
- [444] P.-Y. Chen, H. Zhang, Y. Sharma, J. Yi, and C.-J. Hsieh, “ZOO: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models,” in *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, ACM, pp. 15–26, 2017.
- [445] W. Brendel, J. Rauber, and M. Bethge, “Decision-based adversarial attacks: Reliable attacks against black-box machine learning models,” *arXiv preprint arXiv:1712.04248*, 2017.
- [446] F. Tramèr, A. Kurakin, N. Papernot, I. J. Goodfellow, D. Boneh, and P. D. McDaniel, “Ensemble adversarial training: Attacks and defenses,” in *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 – May 3, 2018, Conference Track Proceedings*, 2018.
- [447] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, “Practical black-box attacks against machine learning,” in *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, ACM, pp. 506–519, 2017.
- [448] A. Athalye, N. Carlini, and D. Wagner, “Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples,” *ICML*, 2018.
- [449] A. Shafahi, M. Najibi, A. Ghiasi, Z. Xu, J. Dickerson, C. Studer, L. S. Davis, G. Taylor, and T. Goldstein, “Adversarial training for free,” *NeurIPS*, 2019.
- [450] C. Xie, Y. Wu, L. van der Maaten, A. Yuille, and K. He, “Feature denoising for improving adversarial robustness,” *CVPR*, 2019.

- [451] Y. Sharma and P.-Y. Chen, “Attacking the Madry defense model with L_1 -based adversarial examples,” *arXiv preprint arXiv:1710.10733*, 2017.
- [452] F. Tramèr and D. Boneh, “Adversarial training and robustness for multiple perturbations,” *arXiv preprint arXiv:1904.13000*, 2019.
- [453] F. Tramèr, J. Behrmann, N. Carlini, N. Papernot, and J.-H. Jacobsen, “Fundamental tradeoffs between invariance and sensitivity to adversarial perturbations,” in *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13–18 July 2020, Virtual Event*, Proceedings of Machine Learning Research, vol. 119, pp. 9561–9571, PMLR, 2020. [Online]. Available: <http://proceedings.mlr.press/v119/tramer20a.html>.
- [454] B. Chen, W. Carvalho, N. Baracaldo, H. Ludwig, B. Edwards, T. Lee, I. Molloy, and B. Srivastava, “Detecting backdoor attacks on deep neural networks by activation clustering,” *arXiv preprint arXiv:1811.03728*, 2018.
- [455] Z. Sun, P. Kairouz, A. T. Suresh, and H. B. McMahan, “Can you really backdoor federated learning?” *arXiv preprint arXiv:1911.07963*, 2019.
- [456] H. Wang, K. Sreenivasan, S. Rajput, H. Vishwakarma, S. Agarwal, J.-y. Sohn, K. Lee, and D. Papailiopoulos, “Attack of the tails: Yes, you really can backdoor federated learning,” *arXiv preprint arXiv:2007.05084*, 2020.
- [457] Y. Ma, X. Zhu, and J. Hsu, “Data poisoning against differentially-private learners: Attacks and defenses,” in *International Joint Conference on Artificial Intelligence (IJCAI), Macao, China*, 2019. [Online]. Available: <https://arxiv.org/abs/1903.09860>.
- [458] R. C. Geyer, T. Klein, and M. Nabi, “Differentially private federated learning: A client level perspective,” *CoRR*, vol. abs/1712.07557, 2017. [Online]. Available: <http://arxiv.org/abs/1712.07557>.

- [459] M. Lécuyer, V. Atlidakis, R. Geambasu, D. Hsu, and S. Jana, “Certified robustness to adversarial examples with differential privacy,” in *2019 IEEE Symposium on Security and Privacy, SP 2019, San Francisco, CA, USA, May 19–23, 2019*, pp. 656–672, 2019. DOI: [10.1109/SP.2019.00044](https://doi.org/10.1109/SP.2019.00044).
- [460] K. Srinathan and C. P. Rangan, “Efficient asynchronous secure multiparty distributed computation,” in *International Conference on Cryptology in India*, Springer, pp. 117–129, 2000.
- [461] V. Mnih and G. E. Hinton, “Learning to label aerial images from noisy data,” in *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pp. 567–574, 2012.
- [462] N. Natarajan, I. S. Dhillon, P. K. Ravikumar, and A. Tewari, “Learning with noisy labels,” in *Advances in Neural Information Processing Systems*, pp. 1196–1204, 2013.
- [463] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel, “Fairness through awareness,” in *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ACM, pp. 214–226, 2012.
- [464] S. Barocas, M. Hardt, and A. Narayanan, *Fairness and Machine Learning*. fairmlbook.org, 2019.
- [465] S. Mitchell, E. Potash, and S. Barocas, “Prediction-based decisions and fairness: A catalogue of choices, assumptions, and definitions,” *arXiv preprint arXiv:1811.07867*, 2018.
- [466] M. J. Kusner, J. Loftus, C. Russell, and R. Silva, “Counterfactual fairness,” in *Advances in Neural Information Processing Systems*, pp. 4066–4076, 2017.
- [467] T. Kamishima, S. Akaho, and J. Sakuma, “Fairness-aware learning through regularization approach,” in *2011 IEEE 11th International Conference on Data Mining Workshops*, IEEE, pp. 643–650, 2011.
- [468] J. Buolamwini and T. Gebru, “Gender shades: Intersectional accuracy disparities in commercial gender classification,” in *Conference on Fairness, Accountability and Transparency*, pp. 77–91, 2018.
- [469] T. Li, M. Sanjabi, and V. Smith, “Fair resource allocation in federated learning,” *arXiv preprint arXiv:1905.10497*, 2019.

- [470] L. Eckhouse, K. Lum, C. Conti-Cook, and J. Ciccolini, “Layers of bias: A unified approach for understanding problems with risk assessment,” *Criminal Justice and Behavior*, vol. 46, no. 2, pp. 185–209, 2019.
- [471] R. Richardson, J. Schultz, and K. Crawford, “Dirty data, bad predictions: How civil rights violations impact police data, predictive policing systems, and justice,” *New York University Law Review Online*, *Forthcoming*, 2019.
- [472] N. Sambasivan, G. Checkley, A. Batool, N. Ahmed, D. Nemer, L. S. Gaytán-Lugo, T. Matthews, S. Consolvo, and E. Churchill, “‘privacy is not for me, it’s for those rich women’: Performative privacy practices on mobile phones by women in South Asia,” in *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*, pp. 127–142, 2018.
- [473] P. Kairouz, J. Liao, C. Huang, and L. Sankar, “Censored and fair universal representations using generative adversarial models,” *arXiv preprint arXiv:1910.00411*, 2020.
- [474] T. Hashimoto, M. Srivastava, H. Namkoong, and P. Liang, “Fairness without demographics in repeated loss minimization,” in *International Conference on Machine Learning*, pp. 1934–1943, 2018.
- [475] Ú. Hébert-Johnson, M. Kim, O. Reingold, and G. Rothblum, “Multicalibration: Calibration for the (computationally-identifiable) masses,” in *International Conference on Machine Learning*, pp. 1944–1953, 2018.
- [476] R. Cummings, V. Gupta, D. Kimpara, and J. Morgenstern, “On the compatibility of privacy and fairness,” in *Proceedings of Fairness in User Modeling, Adaptation and Personalization*, FairUMAP, 2019.
- [477] M. Jagielski, M. J. Kearns, J. Mao, A. Oprea, A. Roth, S. Sharif-Malvajerdi, and J. Ullman, “Differentially private fair learning,” *CoRR*, vol. abs/1812.02696, 2018. [Online]. Available: <http://arxiv.org/abs/1812.02696>.

- [478] E. Bagdasaryan and V. Shmatikov, “Differential privacy has disparate impact on model accuracy,” *CoRR*, vol. abs/1905.12101, 2019. arXiv: [1905.12101](https://arxiv.org/abs/1905.12101). [Online]. Available: <http://arxiv.org/abs/1905.12101>.
- [479] S. Kuppam, R. McKenna, D. Pujol, M. Hay, A. Machanavajjhala, and G. Miklau, “Fair decision making using privacy-protected data,” *CoRR*, vol. abs/1905.12744, 2019. [Online]. Available: <http://arxiv.org/abs/1905.12744>.
- [480] L. Song, R. Shokri, and P. Mittal, “Privacy risks of securing machine learning models against adversarial examples,” in *Proceedings of the ACM Conference on Computer and Communication Security (CCS)*, 2019.
- [481] M. Bertrán, N. Martínez, A. Papadaki, Q. Qiu, M. R. D. Rodrigues, G. Reeves, and G. Sapiro, “Learning adversarially fair and transferable representations,” in *ICML*, 2019.
- [482] C. Feutry, P. Piantanida, Y. Bengio, and P. Duhamel, “Learning anonymized representations with adversarial neural networks,” *CoRR*, vol. abs/1802.09386, 2018. [Online]. Available: <http://arxiv.org/abs/1802.09386>.
- [483] D. Madras, E. Creager, T. Pitassi, and R. Zemel, “Learning adversarially fair and transferable representations,” in *ICML*, 2018.
- [484] B. M. L. Srivastava, A. Bellet, M. Tommasi, and E. Vincent, “Privacy-preserving adversarial representation learning in ASR: Reality or illusion?” In *Annual Conference of the International Speech Communication Association (Interspeech)*, 2019.
- [485] T. S. Brisimi, R. Chen, T. Mela, A. Olshevsky, I. C. Paschalidis, and W. Shi, “Federated learning of predictive models from federated electronic health records,” *Int. J. Med. Informat.*, vol. 112, pp. 59–67, 2018.
- [486] K. Chang, N. Balachandar, C. Lam, D. Yi, J. Brown, A. Beers, B. Rosen, D. L. Rubin, and J. Kalpathy-Cramer, “Distributed deep learning networks among institutions for medical imaging,” *JAMIA*, vol. 25, no. 8, pp. 945–954, 2018.

- [487] A. R. Martin, M. Kanai, Y. Kamatani, Y. Okada, B. M. Neale, and M. J. Daly, “Current clinical use of polygenic scores will risk exacerbating health disparities,” *BioRxiv*, p. 441261, 2019.
- [488] Y. Laguel, K. Pillutla, J. Malick, and Z. Harchaoui, “Device heterogeneity in federated learning: A superquantile approach,” *arXiv preprint arXiv:2002.11223*, 2020.
- [489] C. T. Dinh, N. H. Tran, and T. D. Nguyen, “Personalized Federated Learning with Moreau Envelopes,” in *NeurIPS*, 2020.
- [490] P. Awasthi, C. Cortes, Y. Mansour, and M. Mohri, “Beyond individual and group fairness,” *CoRR*, vol. abs/2008.09490, 2020.
- [491] M. Hardt, E. Price, and N. Srebro, “Equality of opportunity in supervised learning,” in *Advances in Neural Information Processing Systems*, pp. 3323–3331, 2016.
- [492] M. B. Zafar, I. Valera, M. G. Rodriguez, and K. P. Gummadi, “Fairness constraints: Mechanisms for fair classification,” in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 2017.
- [493] D. Ramage and S. Mazzocchi, *Federated analytics: Collaborative data science without data collection*, [Online]. Available: <https://ai.googleblog.com/2020/05/federated-analytics-collaborative-data.html>. Google AI Blog, May 2020.
- [494] R. J. A. Little, “Post-stratification: A modeler’s perspective,” *J. Am. Stat. Assoc.*, vol. 88, no. 423, pp. 1001–1012, 1993.
- [495] D. J. Beutel, T. Topal, A. Mathur, X. Qiu, T. Parcollet, and N. D. Lane, “Flower: A friendly federated learning research framework,” *arXiv preprint arXiv:2007.14390*, 2020.

- [496] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32, pp. 8024–8035, Curran Associates, Inc., 2019. [Online]. Available: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [497] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, *JAX: Composable transformations of Python+NumPy programs*, 2018. [Online]. Available: <http://github.com/google/jax>.
- [498] F. Seide and A. Agarwal, “CNTK: Microsoft’s open-source deep-learning toolkit,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’16, p. 2135, New York, NY, USA: Association for Computing Machinery, 2016. DOI: [10.1145/2939672.2945397](https://doi.org/10.1145/2939672.2945397).
- [499] W. Zhu, P. Kairouz, H. Sun, B. McMahan, and W. Li, “Federated heavy hitters discovery with differential privacy,” *arXiv preprint arXiv:1902.08534*, 2019.
- [500] C. He, S. Li, J. So, X. Zeng, M. Zhang, H. Wang, X. Wang, P. Vepakomma, A. Singh, H. Qiu, X. Zhu, J. Wang, L. Shen, P. Zhao, Y. Kang, Y. Liu, R. Raskar, Q. Yang, M. Annavaram, and S. Avestimehr, “FedML: A research library and benchmark for federated machine learning,” *arXiv preprint arXiv:2007.13518*, 2020.
- [501] N. Rodríguez-Barroso, G. Stipcich, D. Jiménez-López, J. A. Ruiz-Millán, E. Martínez-Cámara, G. González-Seco, M. V. Luzón, M. A. Veganzones, and F. Herrera, “Federated learning and differential privacy: Software tools analysis, the sherpa.ai FL framework and methodological guidelines for preserving data privacy,” *Inform. Fusion*, vol. 64, pp. 270–292, 2020.

- [502] PyVertical Authors, *PyVertical*, 2020. [Online]. Available: <https://github.com/OpenMined/PyVertical>.
- [503] The PaddlePaddle Authors, *PaddlePaddle*, 2019. [Online]. Available: <http://www.paddlepaddle.org/>.
- [504] The Fedlearner Authors, *Fedlearner*, 2020. [Online]. Available: <https://github.com/bytedance/fedlearner>.
- [505] G. Cohen, S. Afshar, J. Tapson, and A. van Schaik, “EMNIST: An extension of MNIST to handwritten letters,” *arXiv preprint arXiv:1702.05373*, 2017.
- [506] S. Caldas, P. Wu, T. Li, J. Konecný, H. B. McMahan, V. Smith, and A. Talwalkar, “LEAF: A benchmark for federated settings,” *arXiv preprint arXiv:1812.01097*, 2018.
- [507] The Google-Landmark-v2 Authors, *Google landmark dataset v2*, 2019. [Online]. Available: <https://github.com/cvdfoundation/google-landmark>.
- [508] J. Luo, X. Wu, Y. Luo, A. Huang, Y. Huang, Y. Liu, and Q. Yang, “Real-world image datasets for federated learning,” *arXiv preprint arXiv:1910.11089*, 2019.
- [509] G. J. Annas, “HIPAA regulations – a new era of medical-record privacy?” *NEJM*, vol. 348, no. 15, pp. 1486–1490, 2003.