

Lecture Notes in Artificial Intelligence

12500

Subseries of Lecture Notes in Computer Science

Series Editors

Randy Goebel

University of Alberta, Edmonton, Canada

Yuzuru Tanaka

Hokkaido University, Sapporo, Japan

Wolfgang Wahlster

DFKI and Saarland University, Saarbrücken, Germany

Founding Editor

Jörg Siekmann

DFKI and Saarland University, Saarbrücken, Germany

More information about this subseries at <http://www.springer.com/series/1244>

Qiang Yang · Lixin Fan · Han Yu (Eds.)

Federated Learning

Privacy and Incentive

Editors

Qiang Yang
WeBank
Shenzhen, China

Lixin Fan
WeBank
Shenzhen, China

Hong Kong University of Science
and Technology
Hong Kong, Hong Kong

Han Yu 
Nanyang Technological University
Singapore, Singapore

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Artificial Intelligence
ISBN 978-3-030-63075-1 ISBN 978-3-030-63076-8 (eBook)
<https://doi.org/10.1007/978-3-030-63076-8>

LNCS Sublibrary: SL7 – Artificial Intelligence

© Springer Nature Switzerland AG 2020

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

Machine learning (ML) has shown significant potential for revolutionizing many important applications such as fraud detection in finance, medical diagnosis in healthcare, or speech recognition in automatic customer service. The traditional approach of training ML models requires large-scale datasets. However, with rising public concerns for data privacy protection, such an approach is facing tremendous challenges. Trust establishment techniques such as blockchains can help users ascertain the origin of the data and audit their usage. Nevertheless, we still require a viable approach to extract value from such trustworthy data and fairly distribute such values to promote collaboration.

Federated learning (FL) is an emerging ML paradigm that aims to help the field of ML adapt to and thrive under the new normal of heightened data privacy concerns and distributively owned data silos. It offers a promising alternative to enable multiple participants to train a globally shared model by exchanging model information without exposing private data.

The protection of data *privacy* is often mandated by the regulatory requirements (e.g., GDPR) in business-to-consumer scenarios. Violations of such regulations can incur hefty fines amounting to the billions. Moreover, in business-to-business settings, participants from the same business sectors may be competitors. This poses further unique challenges for the design of federated *incentives* to fairly account for their contributions and sustain collaboration in the presence of competition. Research works pertaining data privacy protection and incentive mechanism design under FL settings are crucial for the formation and healthy development of FL ecosystems. This is what makes FL unique compared to existing distributed ML paradigms. Therefore, our book focuses on these two main themes.

Although FL training processes are decentralized, without exposing private data, one crux of data privacy protection is to avoid the shared model parameters being exploited by potential adversaries. In this book, we have collected multiple studies on privacy-preserving ML to show the readers potential approaches that can strengthen the privacy aspect of FL.

Despite a wealth of literature on incentive mechanism design exists, the unique settings and challenges facing FL requires meaningful extensions to these approaches. In this book, we have gathered multiple studies on motivating participants to join FL training through rewards (monetary or otherwise) in order to build a sustainable FL ecosystem.

Knowing the theories and techniques about privacy preservation and incentivization under FL is one thing, but successfully applying them in practice also requires non-trivial effort. In this book, we have also included a number of studies on the application of FL in important fields such as recommendation systems and banking.

This book consists of 19 chapters, each of which is a single-blind peer-reviewed submission. Most of the chapters are extensions from workshop or conference

contributions. By providing a well-balanced collection of recent works on *privacy*, *incentive* and the *applications* of FL, the book can help readers gain a more nuanced understanding on how to build a robust and sustainable FL ecosystem and translate the research outcomes into real-world impact. The book is therefore expected to be useful for academic researchers, FL system developers as well as people interested in advanced artificial intelligence topics.

Last but not least, we would like to express our gratitude towards our amazing colleagues, specially Lanlan Chang and Jian Li from the Springer team. Without their help, the publication of this book would not be possible.

September 2020

Qiang Yang
Lixin Fan
Han Yu

Organization

Editorial Team

Qiang Yang	WeBank, China, and The Hong Kong University of Science and Technology, Hong Kong
Lixin Fan	WeBank, China
Han Yu	Nanyang Technological University, Singapore

Program Committee

Kam Woh Ng	WeBank, China
Yilun Jin	The Hong Kong University of Science and Technology, Hong Kong
Tianyu Zhang	WeBank, China

Contents

Privacy

Threats to Federated Learning.	3
<i>Lingjuan Lyu, Han Yu, Jun Zhao, and Qiang Yang</i>	
Deep Leakage from Gradients.	17
<i>Ligeng Zhu and Song Han</i>	
Rethinking Privacy Preserving Deep Learning: How to Evaluate and Thwart Privacy Attacks	32
<i>Lixin Fan, Kam Woh Ng, Ce Ju, Tianyu Zhang, Chang Liu, Chee Seng Chan, and Qiang Yang</i>	
Task-Agnostic Privacy-Preserving Representation Learning via Federated Learning	51
<i>Ang Li, Huanrui Yang, and Yiran Chen</i>	
Large-Scale Kernel Method for Vertical Federated Learning.	66
<i>Zhiyuan Dang, Bin Gu, and Heng Huang</i>	
Towards Byzantine-Resilient Federated Learning via Group-Wise Robust Aggregation.	81
<i>Lei Yu and Lingfei Wu</i>	
Federated Soft Gradient Boosting Machine for Streaming Data.	93
<i>Ji Feng, Yi-Xuan Xu, Yong-Gang Wang, and Yuan Jiang</i>	
Dealing with Label Quality Disparity in Federated Learning.	108
<i>Yiqiang Chen, Xiaodong Yang, Xin Qin, Han Yu, Piu Chan, and Zhiqi Shen</i>	

Incentive

FedCoin: A Peer-to-Peer Payment System for Federated Learning.	125
<i>Yuan Liu, Zhengpeng Ai, Shuai Sun, Shuangfeng Zhang, Zelei Liu, and Han Yu</i>	
Efficient and Fair Data Valuation for Horizontal Federated Learning	139
<i>Shuyue Wei, Yongxin Tong, Zimu Zhou, and Tianshu Song</i>	
A Principled Approach to Data Valuation for Federated Learning	153
<i>Tianhao Wang, Johannes Rausch, Ce Zhang, Ruoxi Jia, and Dawn Song</i>	

A Gamified Research Tool for Incentive Mechanism Design in Federated Learning	168
<i>Zichen Chen, Zelei Liu, Kang Loon Ng, Han Yu, Yang Liu, and Qiang Yang</i>	
Budget-Bounded Incentives for Federated Learning	176
<i>Adam Richardson, Aris Filos-Ratsikas, and Boi Faltings</i>	
Collaborative Fairness in Federated Learning	189
<i>Lingjuan Lyu, Xinyi Xu, Qian Wang, and Han Yu</i>	
A Game-Theoretic Framework for Incentive Mechanism Design in Federated Learning	205
<i>Mingshu Cong, Han Yu, Xi Weng, and Siu Ming Yiu</i>	
Applications	
Federated Recommendation Systems	225
<i>Liu Yang, Ben Tan, Vincent W. Zheng, Kai Chen, and Qiang Yang</i>	
Federated Learning for Open Banking	240
<i>Guodong Long, Yue Tan, Jing Jiang, and Chengqi Zhang</i>	
Building ICU In-hospital Mortality Prediction Model with Federated Learning	255
<i>Trung Kien Dang, Kwan Chet Tan, Mark Choo, Nicholas Lim, Jianshu Weng, and Mengling Feng</i>	
Privacy-Preserving Stacking with Application to Cross-organizational Diabetes Prediction	269
<i>Xiawei Guo, Quanming Yao, James Kwok, Weiwei Tu, Yuqiang Chen, Wenyuan Dai, and Qiang Yang</i>	
Author Index	285