



Delft University of Technology

DistributedML 2022

Chairs' Welcome Message

Laskaridis, Stefanos; Almeida, Mario; Crowcroft, Jon; Li, Ang; Chen, Lydia Y.

Publication date
2022

Document Version
Final published version

Published in
DistributedML 2022 - Proceedings of the 3rd International Workshop on Distributed Machine Learning, Part of CoNEXT 2022

Citation (APA)

Laskaridis, S., Almeida, M., Crowcroft, J., Li, A., & Chen, L. Y. (2022). DistributedML 2022: Chairs' Welcome Message. *DistributedML 2022 - Proceedings of the 3rd International Workshop on Distributed Machine Learning, Part of CoNEXT 2022*, II.

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

December 9, 2022
Rome, Italy



Association for
Computing Machinery

*Advancing Computing
as a Science & Profession*



DistributedML '22

Proceedings of the 3rd International Workshop on
Distributed Machine Learning

Sponsored by:

ACM SIGCOMM

Part of:

CoNEXT'22



**Association for
Computing Machinery**

Advancing Computing as a Science & Profession

The Association for Computing Machinery

**2 Penn Plaza, Suite 701
New York, New York 10121-0701**

Copyright © 2022 by the Association for Computing Machinery, Inc. (ACM). Permission to make digital or hard copies of portions of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyright for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permission to republish from permissions@acm.org or Fax +1 212 869-0481.

For other copying of articles that carry a code at the bottom of the first or last page, copying is permitted provided that the per-copy fee indicated in the code is paid through www.copyright.com.

Notice to Past Authors of ACM-Published Articles

ACM intends to create a complete electronic archive of all articles and/or other material previously published by ACM. If you have written a work that has been previously published by ACM in any journal or conference proceedings prior to 1978, or any SIG Newsletter at any time, and you do NOT want this work to appear in the ACM Digital Library, please inform permissions@acm.org, stating the title of the work, the author(s), and where and when published.

ISBN: 978-1-4503-9922-7

Additional copies may be ordered prepaid from:

**ACM Order Department
PO Box 30777
New York, NY 10087-0777, USA**

Phone: +1 800 342-6626 (USA and Canada)

+1 212 626-0500 (Global)

Fax: +1 212 944-1318

Email: acmhelp@acm.org

Hours of Operation: 8:30 am–4:30 pm ET

Contents

Welcome message and list of PC	ii
 Session 1: Transient resources	
Exploring Learning Rate Scaling Rules for Distributed ML Training on Transient Resources	1
Joel Andre, Foteini Strati, Ana Klimovic (<i>ETH Zurich</i>)	
SmartNICs at Edge for Transient Compute Elasticity	9
Diman Zad Tootaghaj, Anu Mercian (<i>Hewlett Packard Labs</i>); Vivek Adarsh (<i>University of California, Santa Barbara</i>); Mehrnaz Sharifian (<i>University of California Davis</i>); Puneet Sharma (<i>Hewlett Packard Labs</i>)	
 Session 2: Federated Learning	
Rethinking Normalization Methods in Federated Learning	16
Zhixu Du, Jingwei Sun, Ang Li (<i>Duke University</i>); Pin-Yu Chen (<i>IBM Research AI</i>); Jianyi Zhang, Hai "Helen" Li, Yiran Chen (<i>Duke University</i>)	
Towards Predicting Client Benefit and Contribution in Federated Learning From Data Imbalance .	23
Christoph Düsing, Philipp Cimiano (<i>CITEC(Cluster of Excellence Cognitive Interaction Technology), University Bielefeld</i>)	
 Session 3: Efficient & robust ML deployment	
Activation Sparsity and Dynamic Pruning for Split Computing in Edge AI	30
Janek Haberer (<i>Kiel University, Germany</i>); Olaf Landsiedel (<i>Kiel University, Germany and Chalmers University of Technology, Sweden</i>)	
Pelta: Shielding Transformers to Mitigate Evasion Attacks in Federated Learning	37
Simon Queyrut (<i>University of Neuchatel</i>); Yérom-David BROMBERG (<i>Univ. Rennes - Inria - IRISA</i>); Valerio Schiavoni (<i>University of Neuchatel</i>)	
 Author index	 44

DistributedML 2022: Chairs' Welcome Message

Following up the prior two successful versions of DistributedML, it is our great honour and pleasure to welcome you again, this time physically in the 3rd edition of the Distributed Machine Learning Workshop (DistributedML '22). The workshop is co-located with the 18th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '22) and held in Rome, Italy, on the 9th of December 2022.

Distributed ML is a rapidly evolving, interdisciplinary field bringing together techniques from distributed systems, networks and machine learning. The advent of deep learning and the pursuit of artificial intelligence has created previously unfathomable deployments both in the datacenter and in the wild, enabling use-cases ranging from virtual reality and intelligent assistants to robots and self-driving cars. However, be it in the form of training or inference, these workloads pose several challenges pushing the compute, memory, sensing and networking capabilities of devices. As such, ML deployments become innately distributed, as a means to scale out local capabilities by leveraging remote or ambient resources in a distributed fashion in a collaborative manner. Simultaneously, the strive for user-privacy and sustainability gives birth to new training solutions, such as Federated Learning. However, adversarial actors may pose challenges to the robustness of the such deployments.

Be it deployed in the cloud, edge or consumer end, Distributed ML challenges and pushes the limits of today's compute, networking and algorithmic frontier. To this end, the DistributedML workshop comes in a timely manner to provide a forum for ideas coming from different disciplines to be aired so as to solve current challenges and shape the technology of the future.

We would like to thank the ACM CoNEXT organising committee (Giuseppe Bianchi and Alessandro Mei), the workshop co-chairs (Radia Perlman and Paolo Casari) and publication chair (Danilo Giordano) for their guidance throughout the process. We would also like to take the time and thank the DistributedML organising and steering committees, as well as the excellent Technical Program Committee (TPC), with their thorough reviews in a short period of time. Lastly, this workshop would not be possible without the authors of accepted papers and our keynote speakers and panelists.

We look forward to welcoming this year's participants for an exciting and fruitful discussion in the hot topic of Distributed Machine Learning.

The DistributedML organising committee

Stefanos Laskaridis

DistributedML General Chair
Samsung AI Center Cambridge

Mario Almeida

DistributedML General Chair
rain.us

Jon Crowcroft

DistributedML General Chair
University of Cambridge and Alan
Turing Institute

Ang Li

DistributedML'22 Organising Chair
Qualcomm AI

Lydia Y. Chen

DistributedML'22 Organising Chair
TUDelft

ACM DistributedML 2022 Technical Program Committee

Workshop Chairs:

Stefanos Laskaridis, Samsung AI Center Cambridge
Mario Almeida, rain.us
Jon Crowcroft, University of Cambridge & Alan Turing Inst.
Ang Li, Qualcomm AI
Lydia Y. Chen, TU Delft

Website, Publicity, Publication Chairs:

Stefanos Laskaridis, Samsung AI Center Cambridge
Mario Almeida, rain.us

Technical Program Committee:

Jose Cano, Glasgow University
Domenico Ciuonzo, University of Naples “Federico II”
Mallesham Dasari, CMU
Xenofon Foukas, Microsoft Azure for Operators
Samuel Horvath, MBZUAI
Kleomenis Katevas, Telefonica
Hyeji Kim, University of Texas at Austin
Nicolas Kourtellis, Telefonica Research
Anastasios Kyrillidis, Rice University
Fan Lai, University of Michigan
Shen Li, Meta
Mohammad Malekzadeh, Nokia Bell Labs
Anna Maria Mandalari, Imperial College London
Akhil Mathur, Nokia Bell Labs
Mirco Musolesi, University College London
Peter Pietzuch, Imperial College London
Valentin Radu, University of Sheffield
Stojan Trajanovski, Microsoft
Stylianos Venieris, Samsung AI Center Cambridge
Mengwei Xu, BUPT