

2020 IEEE/ACM Fourth Workshop on Deep Learning on Supercomputers (DLS 2020)

**Atlanta, Georgia, USA
11 November 2020**



**IEEE Catalog Number: CFP20W52-POD
ISBN: 978-1-6654-2246-8**

**Copyright © 2020 by the Institute of Electrical and Electronics Engineers, Inc.
All Rights Reserved**

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

****** This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP20W52-POD
ISBN (Print-On-Demand):	978-1-6654-2246-8
ISBN (Online):	978-1-6654-2245-1

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

CURRAN ASSOCIATES INC.
proceedings
.com

2020 IEEE/ACM Fourth Workshop on Deep Learning on Supercomputers (DLS) **DLS 2020**

Table of Contents

Message from the Workshop Chairs .v.....
Workshop Organization .vi.....

Algorithm, Modeling, and Infrastructure

Online-Codistillation Meets LARS; Going Beyond the Limit of Data Parallelism in Deep Learning .1.....
Shogo Murai (Preferred Networks, Inc., Japan), Hiroaki Mikami (Preferred Networks, Inc., Japan), Masanori Koyama (Preferred Networks, Inc., Japan), Shuji Suzuki (Preferred Networks, Inc., Japan), and Takuya Akiba (Preferred Networks, Inc., Japan)

Time-Based Roofline for Deep Learning Performance Analysis .10.....
Yunsong Wang (Lawrence Berkeley National Laboratory), Charlene Yang (Lawrence Berkeley National Laboratory), Steven Farrell (Lawrence Berkeley National Laboratory), Yan Zhang (Lawrence Berkeley National Laboratory), Thorsten Kurth (NVIDIA Corporation), and Samuel Williams (Lawrence Berkeley National Laboratory)

Towards a Scalable and Distributed Infrastructure for Deep Learning Applications .20.....
Bita Hasheminezhad (Louisiana State University, U.S.A.), Shahrzad Shirzad (Louisiana State University, U.S.A.), Nanmiao Wu (Louisiana State University, U.S.A.), Patrick Diehl (Louisiana State University, U.S.A.), Hannes Schulz (Microsoft Research, Canada), and Hartmut Kaiser (Louisiana State University, U.S.A.)

DDL Bench: Towards a Scalable Benchmarking Infrastructure for Distributed Deep Learning .31.....
Matthijs Jansen (SURFsara, University of Amsterdam, The Netherlands), Valeriu Codreanu (SURFsara, The Netherlands), and Ana-Lucia Varbanescu (University of Amsterdam, The Netherlands)

Applications

Vandermonde Wave Function Ansatz for Improved Variational Monte Carlo .40.....
Alberto Acevedo (University of Arizona, USA), Michael Curry (University of Maryland, USA), Shantanu Joshi (University of California, Los Angeles, USA), Brett Leroux (University of California, Davis, USA), and Nicholas Malaya (AMD Research, USA)

TopiQAL: Topic-Aware Question Answering using Scalable Domain-Specific Supercomputers .48...
Hamsa Shwetha Venkataram (Jet Propulsion Laboratory, California Institute of Technology), Chris A Mattmann (Jet Propulsion Laboratory), and Scott Penberthy (Google)

DeepGalaxy: Deducing the Properties of Galaxy Mergers from Images Using Deep Neural Networks .56.....
Maxwell X. Cai (SURF Corporative / Leiden University), Jeroen Bédorf (minds.ai), Vikram A. Saletore (Intel Corporation), Valeriu Codreanu (SURF Corporative), Damian Podareanu (SURF Corporative), Adel Chaibi (Intel Corporation), and Penny X. Qian (Leiden University)

Author Index 63