

MLSys 2020

Austin, Texas

AUSTIN CONVENTION CENTER
500 E Cesar Chavez St, Austin, TX 78701

March 2nd - 4th



BOARD OF DIRECTORS

Ameet Talwalkar (President)
Virginia Smith (Secretary)
Inderjit Dhillon (Treasurer)
Matei Zaharia

STEERING COMMITTEE

Jennifer Chayes
Bill Dally
Jeff Dean
Michael I. Jordan
Yann LeCun
Fei-Fei Li
Alex Smola
Dawn Song
Eric Xing

ORGANIZING COMMITTEE

General Chair:
Inderjit Dhillon

Program Chairs:
Dimitris Papailiopoulos
Vivienne Sze

Workshop Chairs:
Ralf Herbrich
Theodoros Rekatsinas

Publications Chair:
Jakub Konečný

Publicity Chair:
Gauri Joshi

Artifact Evaluation Chairs:
Grigori Fursin
Gennady Pekhimenko

Demo Chairs:
Jing Li
Ce Zhang

Web Chair:
Siddhartha Sen

PROGRAM COMMITTEE

Jayadev Acharya
Dan Alistarh
Mohammad Alizadeh
Gustavo Alonso
David Andersen
Peter Bailis
Michael Carbin
Nicholas Carlini
Bryan Catanzaro
Tianqi Chen
Yu-Hsin Chen
Yuejie Chi
Minsik Cho
Christopher De Sa
Greg Diamos
Alex Dimakis
Lara Dolecek
Pradeep Dubey
Hadi Esmaeilzadeh

Kayvon Fatahalian
Phillip Gibbons
Garth Gibson
Anna Goldie
Justin Gottschlich
Kim Hazelwood
Cho-Jui Hsieh
Prateek Jain
Gauri Joshi
Rania Khalaf
Jakub Konečný
Farinaz Koushanfar
Sanmi Koyejo
Tim Kraska
Arun Kumar
Anastasios Kyrillidis
Aparna Lakshmiratan
Kangwook Lee
Hai (Helen) Li

Jing Li
Siyuan Ma
Diana Marculescu
Erik Meijer
Xiangrui Meng
Ioannis Mitliagkas
Rajat Monga
Andreas Moshovos
Jennifer Myers
Rob Nowak
Gennady Pekhimenko
Kannan Ramchandran
Alex Ratner
Theodoros Rekatsinas
Bitva Rouhani
Hanie Sedghi
Siddhartha Sen
Yakun Sophia Shao

Virginia Smith
Evan Sparks
Suvrit Sra
Martin Takac
Ameet Talwalkar
Lingjia Tang
Joaquin Vanschoren
Shivaram Venkataraman
Marian Verhelst
Markus Weimer
Paul Whatmough
Andrew Wilson
Stephen Wright
Carole-Jean Wu
Neeraja Yadwadkar
Eiko Yoneki
Matei Zaharia
Ce Zhang

MLSys 2020 Sponsors

facebook

Qualcomm

IBM

 databricks

 CISCO



amazon

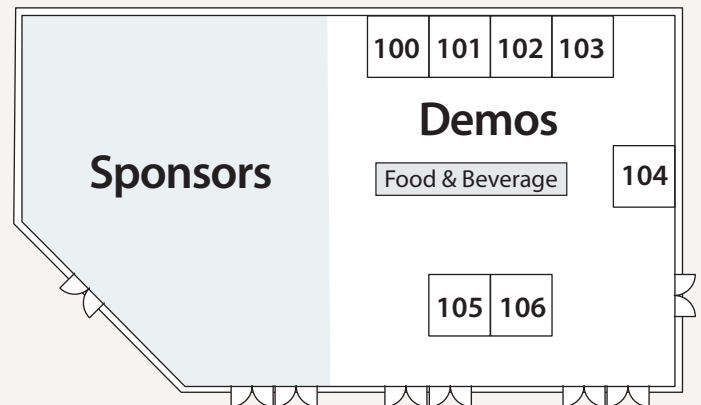
 Weights & Biases

 DeepMind



Demonstrations: Monday 6 – 9 pm

- #100 Fairness Indicators Demo: Scalable Infrastructure for Fair ML Systems
- #101 Demonstration of Ballet: A Framework for Open-Source Collaborative Feature Engineering
- #102 PlaidML EDSL: Combining Programmability and Portability
- #103 Skyline: Interactive In-editor Performance Visualizations and Debugging for DNN Training
- #104 Air Learning: An End To End Learning Gym For Aerial Robots
- #105 Visual Face Re Identification on a Battery-Powered Multi-Core Edge Processor
- #106 3D Adversarial Object against MSF-based Perception in Autonomous Driving



MLSys Agenda: Monday March 2nd

7:00 - 7:45 am Breakfast & Registration

7:45 - 8:00 am Opening Remarks

8:00 - 10:05 am **Session 1: Distributed and Parallel Learning Algorithms**

- A System for Massively Parallel Hyperparameter Tuning
- PLink: Discovering and Exploiting Locality for Accelerated Distributed Training on the public Cloud
- Federated Optimization in Heterogeneous Networks
- BPPSA: Scaling Back-propagation by Parallel Scan Algorithm
- Distributed Hierarchical GPU Parameter Server for Massive Scale Deep Learning Ads Systems

10:05 - 10:30 am Coffee Break

10:30 - 12:10 pm **Session 2: Efficient Model Training**

- Resource Elasticity in Distributed Deep Learning
- SLIDE: Training Deep Neural Networks with Large Outputs on a CPU faster than a V100-GPU
- FLEET: Flexible Efficient Ensemble Training for Heterogeneous Deep Neural Networks
- Breaking the Memory Wall with Optimal Tensor Rematerialization

12:10 - 1:30 pm Lunch Provided

1:30 - 2:30 pm **Keynote: Chris Ré: Theory and Systems for Weak Supervision**

2:30 - 4:00 pm **Session 3: Efficient Inference and Model Serving**

- What is the State of Neural Network Pruning?
- SkyNet: a Hardware-Efficient Method for Object Detection and Tracking on Embedded Systems
- MNN: A Universal and Efficient Inference Engine
- Willump: A Statistically-Aware End-to-end Optimizer for Machine Learning Inference

4:10 - 4:30 pm Coffee Break

4:30 - 6:10 pm **Session 4: Model/Data Quality and Privacy**

- Attention-based Learning for Missing Data Imputation in HoloClean
- Privacy-Preserving Bandits
- Understanding the Downstream Instability of Word Embeddings
- Model Assertions for Monitoring and Improving ML Models

6:10 - 9:00 pm Demonstrations

6:30 - 9:00 pm Posters & Reception (dinner + drinks)

MLSys Agenda: Tuesday March 3rd

7:00 - 8:00 am Breakfast & Registration

8:00 - 10:05 am **Session 5 : ML programming models and abstractions & ML applied to systems**

- AutoPhase: Juggling HLS Phase Orderings in Random Forests with Deep Reinforcement Learning
- Automatically batching control-intensive programs for modern accelerators
- Predictive Precompute with Recurrent Neural Networks
- Sense & Sensitivities: The Path to General-Purpose Algorithmic Differentiation
- Ordering Chaos: Memory-Aware Scheduling of Irregularly Wired Neural Networks for Edge Devices

10:05 - 10:30 am Coffee Break

10:30 - 12:10 pm **Session 6: Efficient Inference and Model Serving**

- Fine-Grained GPU Sharing Primitives for Deep Learning Applications
- Improving the Accuracy, Scalability, and Performance of Graph Neural Networks with Roc
- OPTIMUS: OPTImized matrix MULTiplication Structure for Transformer neural network accelerator
- PoET-BiN: Power Efficient Tiny Binary Neurons

12:10 - 1:30 pm Lunch Provided

1:30 - 2:30 pm **Keynote: Shafi Goldwasser: The Emerging Role of Cryptography in Trustworthy AI**

2:30 - 4:10 pm **Session 7: Quantization of Deep Neural Networks**

- Memory-Driven Mixed Low Precision Quantization for Enabling Deep Network Inference on Microcontrollers
- Trained Quantization Thresholds for Accurate and Efficient Fixed-Point Inference of Deep Neural Networks
- Riptide: Fast End-to-End Binarized Neural Networks
- Searching for Winograd-aware Quantized Networks

4:10 - 4:30 pm Coffee Break

4:30 - 6:10 pm **Session 8: Efficient Model Training 2**

- Blink: Fast and Generic Collectives for Distributed ML
- A Systematic Methodology for Analysis of Deep Learning Hardware and Software Platforms
- MotherNets: Rapid Deep Ensemble Learning
- MLPerf Training Benchmark

6:10 - 6:15 pm Closing Remarks & MLSys 2021

MLSys Workshops: Wednesday March 4th

9:00 am – 5:30 pm

7 am - 8 am Breakfast & Registration

8 am - 5:30 pm Workshops (Check the schedule for each workshop start time)

Benchmarking Machine Learning Workloads on Emerging Hardware

Tom St John · Murali Emani

Level 3, Room 6

Automated Machine Learning For Networks and Distributed Systems

Behnaz Arzani · Bitu Darvish Rouhani

Level 3, Room 10

SARA: Secure and Resilient Autonomy

Pradip Bose · Nandhini Chandramoorthy · Augusto Vega · Karthik Swaminathan

Level 1, Room 3

Resource-Constrained Machine Learning

Yaniv Ben Itzhak · Nina Narodytska · Christopher Aberger

Level 3, Room 8

MLOps Systems

Debo Dutta · Matei Zaharia · Ce Zhang

Level 3, Room 5

Software-Hardware Codesign for Machine Learning Workloads

Ritwik Gupta · John Wohlbiel · Tze Meng Low

Level 3, Room 9

On-Device Intelligence

Vikas Chandra · Pete Warden · Ganesh Venkatesh · Yingyan Lin

Ballroom A

10 am - 10:30 am Coffee Break

12 pm - 2 pm Lunch on your own

3:30 pm - 4 pm Coffee Break