

FastPath 2019 - CALL FOR PAPERS
International Workshop on Performance Analysis of Machine Learning Systems
March 24, 2019 - Madison, Wisconsin, United States
<https://tinyurl.com/2019-FastPath>

In conjunction with ISPASS 2019:
<http://www.ispass.org/ispass2019>

SUMMARY

FastPath 2019 brings together researchers and practitioners involved in cross-stack hardware/software performance analysis, modeling, and evaluation for efficient machine learning systems. Machine learning demands tremendous amount of computing. Current machine learning systems are diverse, including cellphones, high performance computing systems, database systems, self-driving cars, robotics, and in-home appliances. Many machine-learning systems have customized hardware and/or software. The types and components of such systems vary, but a partial list includes traditional CPUs assisted with accelerators (ASICs, FPGAs, GPUs), memory accelerators, I/O accelerators, hybrid systems, converged infrastructure, and IT appliances. Designing efficient machine learning systems poses several challenges.

These include distributed training on big data, hyper-parameter tuning for models, emerging accelerators, fast I/O for random inputs, approximate computing for training and inference, programming models for a diverse machine-learning workloads, high-bandwidth interconnect, efficient mapping of processing logic on hardware, and cross system stack performance optimization. Emerging infrastructure supporting big data analytics, cognitive computing, large-scale machine learning, mobile computing, and internet-of-things, exemplify system designs optimized for machine learning at large.

TOPICS

FastPath seeks to facilitate the exchange of ideas on performance optimization of machine learning/AI systems and seeks papers on a wide range of topics including, but not limited to:

- o Workload characterization, performance modeling and profiling of machine learning applications
- o GPUs, FPGAs, ASIC accelerators
- o Memory, I/O, storage, network accelerators
- o Hardware/software co-design
- o Efficient machine learning algorithms
- o Approximate computing in machine learning
- o Power/Energy and learning acceleration
- o Software, library, and runtime for machine learning systems
- o Workload scheduling and orchestration
- o Machine learning in cloud systems
- o Large-scale machine learning systems
- o Emerging intelligent/cognitive system
- o Converged/integrated infrastructure
- o Machine learning systems for specific domains, e.g., financial, biological, education, commerce, healthcare

SUBMISSION

Prospective authors must submit a 2-4 page extended abstract electronically at:
<https://easychair.org/conferences/?conf=fastpath2019>

Authors of selected abstracts will be invited to give a 30-min presentation at the workshop.

KEY DATES

Submission: February 8, 2019

Notification: February 22, 2019

Final Materials / Workshop: March 24, 2019

ORGANIZERS

General Chair: Erik Altman

Program Committee Chairs: Zehra Sura, Parijat Dube