

Proposal: Workshop at ISC 2020

Machine Learning on HPC Systems (MLHPCS)

Workshop format: ½ day with proceedings

Abstract

Over the last few years, Machine Learning/Deep Learning (ML/DL) has become an important research topic in the High Performance Computing (HPC) community. Bringing new users and data intensive applications on HPC systems, Machine Learning is increasingly affecting the design and operation of compute infrastructures. On the other hand, the Machine Learning community is just getting started to utilize the performance of HPC, leaving many opportunities for better parallelization and scalability. The intent of this workshop is to bring together researchers and practitioners to discuss three key topics in the context of High Performance Computing and Machine Learning/Deep Learning: parallelization and scaling of ML/DL algorithms, HPC system design and optimization for ML/DL workloads and ML/DL applications on HPC systems.

Topics / Scope

The aim of the workshop is to provide a platform for technical discussions, work in progress and the presentation of unsolved problems, which is complementary to the “Machine Learning Day” in the main conference program.

- Unsolved problems in ML/DL on HPC systems
- Scalable Machine Learning / Deep Learning algorithms
- Parallelization techniques
- Libraries for ML/DL
- Tools + workflows for ML/DL on HPC systems
- Optimized HPS System design / setup for efficient ML/DL
- ML Applications on HPC Systems

Organizing committee

- Janis Keuper (IMLA)
 - Contact: keuper@imla.ai
 - Institution 1: Institute for Machine Learning and Analytics (IMLA), Offenburg University, Germany
 - Institution 2: CC-HPC, Fraunhofer ITWM, Kaiserslautern, Germany
 - Job title. Full Professor
 - CV:

Janis Keuper is full professor for Data Science and Analytics at the Institute for Machine Learning and Analytics (IMLA), Offenburg University and scientific advisor at the "Large Scale Machine Learning" group at the Fraunhofer Competence Center for High Performance Computing. His current research is focused on scalable machine learning systems, especially Deep Learning. Before joining IMLA in 2019, he was a Group Leader at Fraunhofer ITWM and the Intel Visual Computing Institute (Saarbrücken, Germany). Janis was the chair of the Deep Learning tracks at the ISC Supercomputing 2017 and 2018 conference and member of the organizing committee of the "Machine Learning in HPC" Workshop at the ACM Supercomputing 2018/2019 conferences.
- Juan J. Durillo (LRZ)
 - Contact: durillo@lrz.de
 - Institution: Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities
 - Jobt title: Senior Researcher
 - CV: Juan J. Durillo received the M.S. and Ph.D. degrees in computer science from the University of Málaga, Spain, in 2006 and 2011, respectively. From 2011 to 2017, he worked as an assistant professor at the University of Innsbruck, Austria. During this period, Dr. Durillo performed research on the intersection of artificial intelligence/machine learning and compilers for optimizing source code. This work led to his Habilitation in Computer Science in 2018. Dr. Durillo has authored more than 50 publications in international journals, conferences and book chapters. As a lecturer, he has taught courses on optimization, operating systems, GPU programming and C++. Since 2018, he works as a Scientist at the BADW-LRZ. His research interests include automatic tuning of scientific applications, multi criteria optimization, GPU computing, machine learning and data analytics.

- Dennis Hoppe (HLRS)
 - Contact: hoppe@hlrs.de
 - Institution: High Performance Computing Centre Stuttgart (HLRS), Stuttgart, Germany
 - Job title: Head of Service Management and Business Processes
 - CV: Dennis Hoppe is heading the department SANE for service management and business processes at HLRS. After he received his master's degree in computer science and media in 2010 from the Bauhaus-University Weimar, Mr. Hoppe was employed as a research assistant at the Bauhaus-University until 2012. In Weimar, he worked on national research projects such as CAIR, where he gained a strong background in information retrieval, data mining, and big data analytics. Before joining HLRS in 2014, Dennis Hoppe worked as a software engineer at ePages, the leading cloud-based e-commerce solution provider in Europe. At HLRS, Dennis Hoppe has been involved in multiple European research projects such as EXCESS (monitoring, increase energy efficiency) and DreamCloud (improve resource allocation). He is currently coordinating the high-performance data analytics project CATALYST at HLRS in cooperation with the HPC manufacturer Cray; goal of the project is the evaluation of the convergence of AI and HPC.

- Jenia Jitsev (JSC)
 - Contact: j.jitsev@fz-juelich.de
 - Institution: Juelich Supercomputing Center (JSC), Research Center Juelich (FZJ)
 - Job title: Head of Cross-Sectional Team Deep Learning (CST-DL)
 - CV: Jenia Jitsev is senior researcher, leading the Cross-Sectional Team Deep Learning (CST-DL) at the Jülich Supercomputing Center (JSC). He obtained Diploma degree in computer science (with minor in psychology) from the University Bonn, with focus on neuroscience and machine learning, and his PhD in computer science at the Frankfurt Institute for Advanced Studies (FIAS) and Goethe University of Frankfurt, working on plasticity and unsupervised learning in recurrent hierarchical neural networks of the visual cortex, with applications to object and face recognition. For his work on reinforcement learning in the basal ganglia, he received Best Paper Award from International Society of Neural Networks and IEEE. Currently, he serves as scientific head of Helmholtz AI Local at JSC. His long-term research agenda aims on establishing closed-loop continual learning in multi-task scenarios using end-to-end trainable pipelines, capable of growing generic models from incoming streams of data, extracting knowledge and skills transferable across different domains and tasks. From a high performance computing (HPC) perspective, his interest is in scaling up training and inference procedure in deep neural networks across multiple GPUs or other accelerators.

- Sunna Torge (ZIH Dresden)
 - Contact: sunna.torge@tu-dresden.de
 - Institution: Center for Information Services and High Performance Computing (ZIH), Technische Universität Dresden
 - Job title: Senior Researcher
 - CV: Sunna Torge is a senior researcher in the national AI and Big Data competence center ScaDS.AI Dresden/Leipzig at ZIH (TU Dresden). She obtained a diploma in mathematics (minor physics) from the University of Freiburg and a PhD in computer science with focus on mathematical logic and automated deduction. After working in research in the man-machine-interface group at Sony International (Europe) and the machine learning group at University of Freiburg she was a full professor for theoretical computer science at the University of Applied Science Furtwangen. After her move to Dresden Sunna Torge worked in machine learning and data analytics groups within the TU Dresden and Fraunhofer Institute for Transportation and Infrastructure Systems with focus on text analysis, and data analytics in the traffic domain. Her main research focus currently is on text and sequence analysis on large data.

Program committee

- Peter Labus (Fraunhofer ITWM)
- Michael Steyer (Intel)
- Jamal Toutouh (MIT Computer Science and Artificial Intelligence Lab)
- René Jäkel (ZIH, TU Dresden)
- Janis Keuper (IMLA)
- Juan J. Durillo (LRZ)
- Dennis Hoppe (HLRS)
- Jenia Jitsev (JSC)
- Sunna Torge (ZIH, TU Dresden)
- ...further PC members pending

Schedule

Plan: ½ day has 3.5 hours without breaks

- 2 invited Talks (a 45min)
- 6-8 contributed talks (a 20min / 15min)

Candidates for Invited Talks (to be confirmed):

- Brian van Esen, LLNL
- Daniel Soudry, Technion