

PROGRAM

Time	Speaker	Title
09:00-9:20	Boris Ginsburg, Intel	Opening Remarks
Hardware Acceleration for Deep Learning		
09:29-9:45	Amir Khosrowshahi, Nervana Systems	Computer Architectures for Deep Learning
9:45-10:10	Eric Chung, Microsoft Research	Accelerating Deep Convolutional Neural Networks Using Specialized Hardware in the Datacenter
10:10-10:35	Vinayak Gokhale, Purdue University	A Hardware Accelerator for Convolutional Neural Networks
10:35-11:00	Paul Burchard Goldman Sachs	Hardware Acceleration for Communication-Intensive Algorithms
11:00-11:30	COFFE	
Machine Learning Workloads Analysis		
11:30-12:00	Jonathan Pearce, Intel Labs	You Have No (Predictive) Power Here, SPEC!
12:00-12:30	Scott Beamer, UC Berkeley,	Graph Processing Bottlenecks
12:30-13:30	LUNCH	
Neuromorphic Engineering		
13:30-14:00	James E. Smith Univ. Wisconsin – Madison	Biologically Plausible Spiking Neural Networks
14:00-14:30	Giacomo Indiveri, Institute of Neuroinformatics, Univ. of Zurich and ETH Zurich	Neuromorphic circuits and for building autonomous cognitive systems
14:30-15:00	Yiran Chen, Univ. of Pittsburgh	Hardware Acceleration for Neuromorphic Computing: An Evolving View
15:00-15:30	Mikko Lipasti, Univ. of Wisconsin – Madison	Mimicking the Self-Organizing Properties of the Visual Cortex
15:30-16:00	COFFE	
Hardware Acceleration for Machine Learning		
16:00-16:30	Shai Fine, Intel	Machine Learning Building Blocks
16:30-17:00	Chunkun Bo, University of Virginia	String Kernel Testing Acceleration using the Micron Automata Processor
17:00-17:30	Ran Ginosar, Technion	Accelerators for Machine Learning of Big Data