



Contribution ID: 102

Type: **Poster plus Minioral**

FastML Science Benchmarks: Accelerating Real-Time Scientific Edge Machine Learning

Applications of machine learning (ML) are growing by the day for many unique and challenging scientific applications. However, a crucial challenge facing these applications is their need for ultra low-latency and on-detector ML capabilities. Given the slowdown in Moore's law and Dennard scaling, coupled with the rapid advances in scientific instrumentation that is resulting in growing data rates, there is a need for ultra-fast ML at the extreme edge. Fast ML at the edge is essential for reducing and filtering scientific data in real-time to accelerate science experimentation and enable more profound insights. To accelerate real-time scientific edge ML hardware and software solutions, we need well-constrained benchmark tasks with enough specification to be generically applicable and accessible. These benchmarks can guide the design of future edge ML hardware for scientific applications capable of meeting the nanosecond and microsecond level latency requirements. To this end, we present an initial set of scientific ML benchmarks, covering a variety of ML and embedded system techniques.

Minioral

No

IEEE Member

No

Are you a student?

No

Authors: HAWKS, Benjamin (Fermi National Accelerator Lab); HERWIG, Christian (Fermi National Accelerator Lab. (US)); DUARTE, Javier Mauricio (Univ. of California San Diego (US)); MUHIZI, Jules (Harvard University); PRAKASH, Shvetank (Harvard University); REDDI, Vijay (Harvard University); TRAN, Nhan (Fermi National Accelerator Lab. (US))

Presenter: HERWIG, Christian (Fermi National Accelerator Lab. (US))

Session Classification: Mini Oral - III