FastML Science Benchmarks: Accelerating Real Time Scientific Edge Machine Learning

Experiments continuously evolve to probe ***shorter distances + timescales***.

- Powerful detectors → large data volumes
- Require some shrewd selection or distillation ***at the edge***.
  - Increasingly with ML, deep neural nets

Benchmarks guide development of next-gen edge ML hardware + software.

- Science tasks eclipse current standards!

We propose a **new set of benchmarks** including representative tasks across a range of scientific domains.

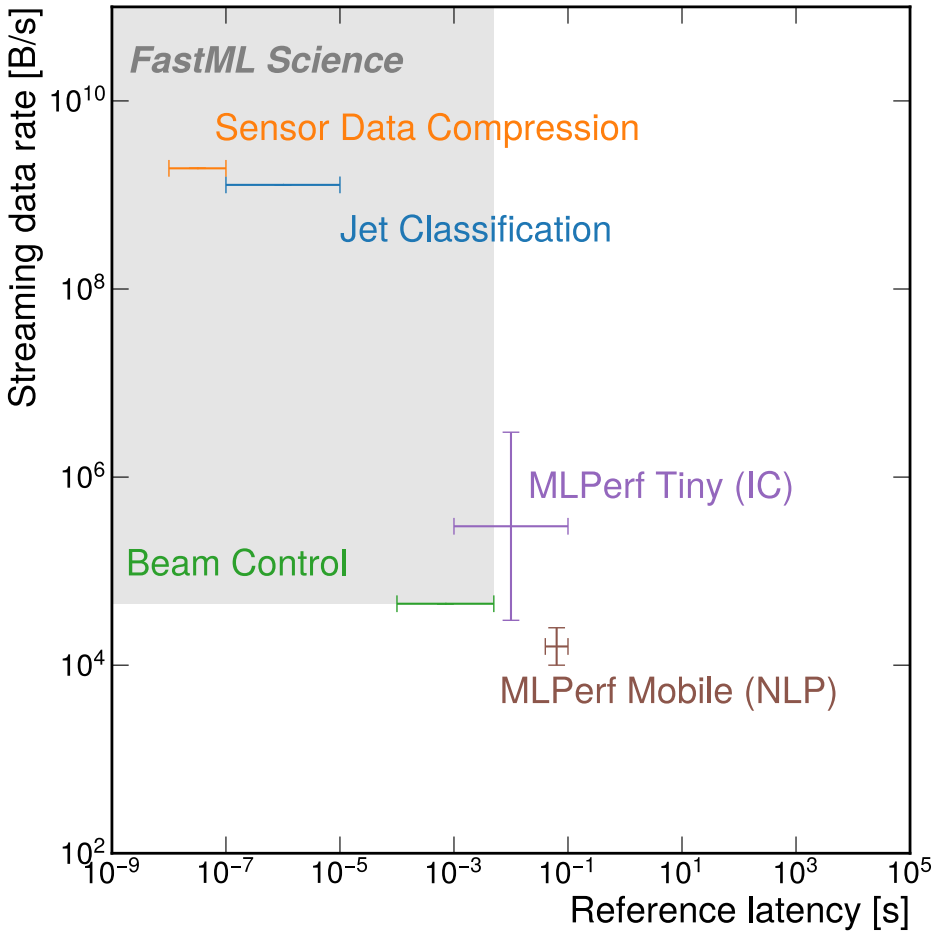

Figure: Reference latencies and streaming input data rates for common benchmarks and those proposed in this work.

FastML Science Benchmarks: Accelerating Real Time Scientific Edge Machine Learning

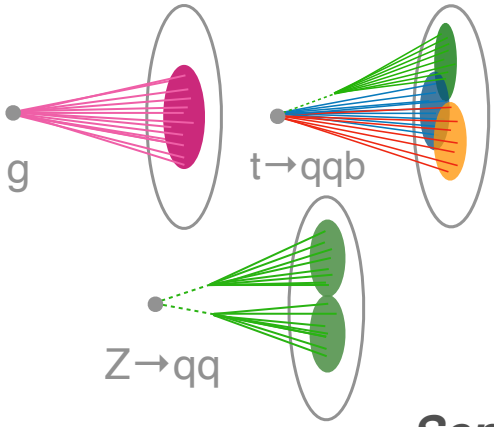

Supervised classification of particle jets

- LHC Trigger: 100 TB/s (Virtex Ultrascale+)
- 1 μs latency, 150 ns pipeline

Sensor data compression

- Next-gen “imaging calorimeter”
- ASIC → area, power constraints

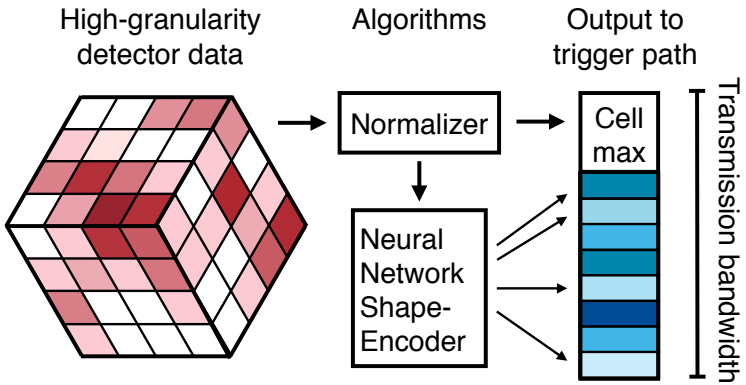

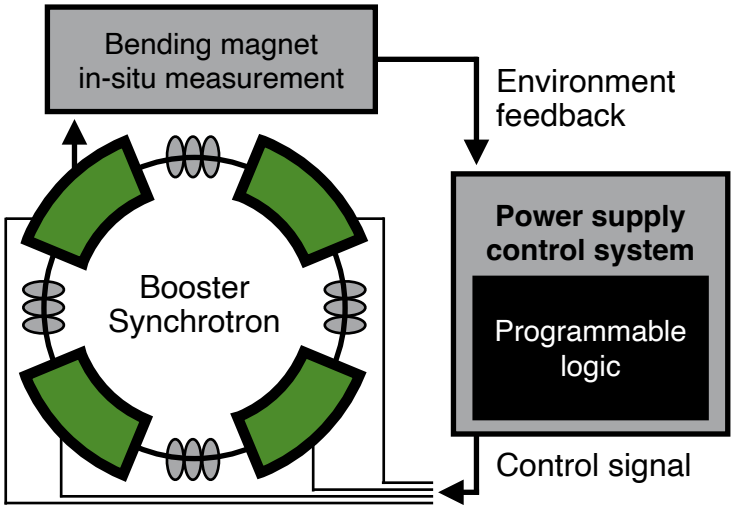

Accelerator beam controls

- Controller (Arria10) interacts with ‘environment’
  - Reinforcement learning
- Inputs from 50 devices across accel. complex.