

tinyMLPerf: Benchmarking Ultra-low Power Machine Learning Systems

Vijay Janapa Reddi, MLPerf Inference Chair, Associate Professor in the John A. Paulson School of Engineering and Applied Sciences, Harvard University

Tiny machine learning (ML) is poised to drive enormous growth within the IoT hardware and software industry. Measuring the performance of these rapidly proliferating systems, and comparing them in a meaningful way presents a considerable challenge; the complexity and dynamicity of the field obscure the measurement of progress and make embedded ML application and system design and deployment intractable. To foster more systematic development, while enabling innovation, a fair, replicable, and robust method of evaluating tinyML systems is required. A reliable and widely accepted tinyML benchmark is needed. To fulfill this need, tinyMLPerf is a community-driven effort to extend the scope of the existing MLPerf benchmark suite (mlperf.org) to include tinyML systems. With the broad support of over 75 member organizations, the tinyMLPerf group has begun the process of creating a benchmarking suite for tinyML systems. The talk presents the goals, objectives, and lessons learned (thus far), and welcomes others to join and contribute to tinyMLPerf.