EECS 598 Section 005
Fall 2018
Computer Hardware Design for Machine Learning

Instructor: Zhengya Zhang
zhengya@umich.edu

Lectures: MW 10:30 – 12:00 pm, 3427 EECS

Prerequisite: EECS 427 or EECS 470

Units: 3 (or 4 with an optional project)

Grading: 40% Topic presentation
40% Review of presentations
20% Term paper

Course Description:

Machine learning has evolved rapidly in the last decade and it has become ubiquitous in applications from smart devices to self-driving cars. A key enabler of modern machine learning is the availability of low-cost, high-performance computer hardware, such as graphics processing units (GPUs) and specialized accelerators such as Google's tensor processing unit (TPU). New machine learning applications constantly impose new requirements and constraints on the hardware design. Hardware implementations must fit increasingly stringent area and power envelope. This course will survey the latest architecture and circuit designs for machine learning applications. Paper reviews and presentation will be the essential parts of this course. An optional unit can be earned by benchmarking or prototyping selected designs that leads to insightful conclusions.

Topics:

Deep (convolutional) neural nets
Spiking neural nets
Neuro-inspired designs
Analog and mixed-signal implementations
Process in memory and applications
Emerging compute substrates: RRAM and MRAM
Alternative applications: speech, stereo vision
Graph processing
Point clouds