NEW COURSE ANNOUNCEMENT FOR FALL 2019

EECS 498/598: Brain-Inspired Computing: Models, Architectures, and Programming

Time: Tuesday and Thursday 1:30 to 3:00 pm
Instructor: Prof. Pinaki Mazumder
Phone: 734-763-2107; e-mail: mazum@umich.edu

Brain-inspired computing is a subset of AI-based machine learning and is generally referred to both deep and shallow artificial neural networks (ANN) and spiking neural networks (SNN). Deep convolutional neural networks (CNN) have made pervasive market inroads in numerous commercial applications and their software implementations are widely studied in computer vision, speech processing and other courses. The purpose of this course will be to study the wide gamut of shallow and deep neural network models, the methodologies for specialized hardware design of popular learning algorithms, as well as adapting hardware architectures on crossbar fabrics of emerging technologies such as memristors and spin torque nanomagnetic devices. Existing software development tools such as TensorFlow, Caffe, and PyTorch will be leveraged to teach various aspects of neuromorphic designs.

Prerequisites: Senior undergrad and grad student standing in Electrical Engineering, Computer Engineering, Computer Science or Applied Physics program.

Outline: i) Fundamentals of brain-inspired computing and history of neural computing, ii) Basics of linear algebra and probability theory needed for modeling of neural networks, iii) Deep learning by convolutional neural networks such as AlexNet, VGG, GoogLeNet, and ResNet, iv) Deep Neural Net (DNN) software development resources, v) Shallow neural networks – Perceptron, Hopfield network, Boltzmann machine, Recurrent neural network, and Kohonen's self-organizing map, vi) Learning models for artificial and spiking neural networks (ANN & SNN) such as spike timing dependent plasticity (STDP), Q-learning, actor-critic reinforcement learning, supervised learning, and back-propagation algorithms, vii) Experimental brain-like large machines such as Caviar, TrueNorth, Neurogrid, and SpiNNaker, viii) Commercial hardware accelerators such as NVDIA's graphics processing unit (GPU), Google's tensor processing unit (TPU), and Intel's vision processing unit (VPU), ix) Application-specific VLSI chips capable of STDP learning, actor/critic reinforcement learning, and Q-learning, and x) Neuromporphic circuits using emerging technologies such as memristors, spin transfer torque devices, and photonic devices.

Evaluation Criteria: A) Critiquing of papers and presentation in class (40%), B) One programming assignment on deep learning (20%) using DNN simulation tool, and C) An end-of-the-term project (40%), which may include i) topical review or survey of related publications, ii) implementation of a learning based application using simulation tools, or iii) developing a neural hardware in ASIC or FPGA.

Teaching materials: The course will mainly rely on archival journal papers along with following reference books:

1. Deep Learning by Aaron C. Courville, Ian Goodfellow, and Yoshua Bengio, MIT Press, 2015.
2. Learning in Energy-Efficient Neuromorphic Computing: Algorithm and Architecture Co-Design by Nan Zheng and Pinaki Mazumder, John Wiley & Sons, USA, 2019.
3. Neuromorphic Circuits for Nanoscale Devices by Pinaki Mazumder, et al., River Publishing, UK, 2019.