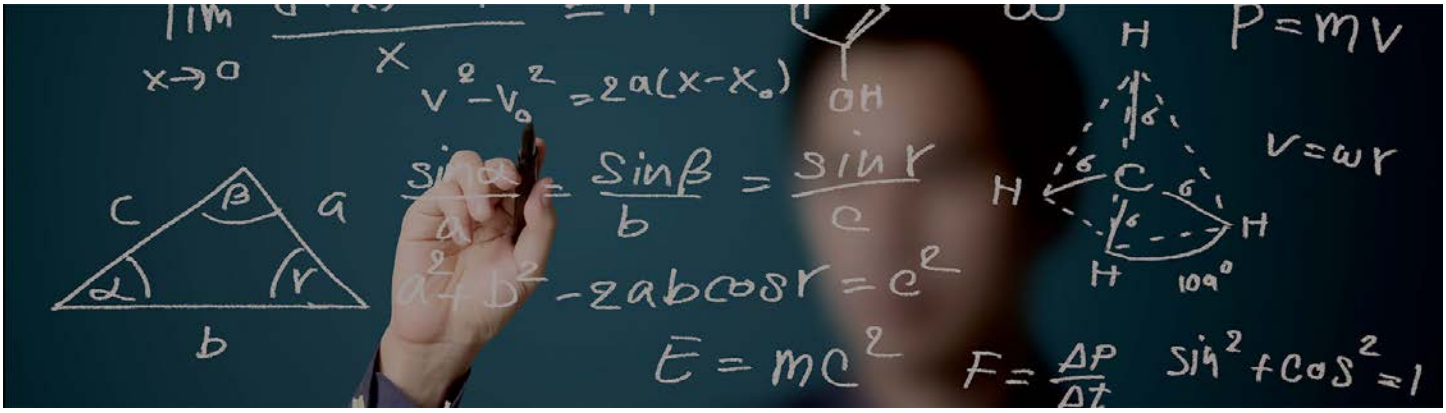

Spring 2019

RENSSELAER POLYTECHNIC INSTITUTE

DEPARTMENT OF MATHEMATICAL SCIENCES COLLOQUIUM

"Deep Learning with Graph Structured Data: Methods, Theory, and Applications"

Abstract: Graphs are universal representations of pairwise relationship. A trending topic in deep learning is to extend the remarkable success of well-established neural network architectures (e.g., CNN and RNN) for Euclidean structured data to irregular domains, including notably, graphs. A proliferation of graph neural networks (e.g., GCN) emerged recently, but the scalability challenge for training and inference persists. The essence of the problem is the prohibitive computational cost of computing a mini-batch, owing to the recursive expansion of neighborhoods. We propose a scalable approach, coined FastGCN, based on neighborhood sampling to reduce the mini-batch computation. FastGCN achieves orders of magnitude improvement in training time, compared with a standard implementation of GCN. Predictions remain comparably accurate. A curious question for this approach is why stochastic gradient descent (SGD) training ever converges. In the second part of this talk, we analyze that the gradient estimator so computed is not unbiased but consistent. We thus extend the standard SGD results for unbiased gradients to consistent gradients and show that their convergence behaviors are similar. These results are important and may spawn new interest in the machine learning community, since in many learning scenarios unbiased estimators may not be efficient to compute, and hence other nonstandard but fast gradient estimators serve as sound alternatives.

Jie Chen (IBM Research)

Monday, March 25, 2019

4-5pm

Amos Eaton 214

Host: Yangyang Xu

Refreshments served 3:30-4pm Amos Eaton 4th Floor Lounge