

Syllabus for IE 598: Neural Networks and Deep Learning

Instructor: Justin Sirignano

TAs: Logan Courtney, Biplab Deka

Course Overview:

This course provides an introduction to neural networks and recent advances in deep learning. Topics include training and implementation of neural networks, convolution neural networks, recurrent neural networks (LSTM and gated recurrent), residual networks, reinforcement learning, and Q-learning with neural networks. A part of the course will especially focus on recent work in deep reinforcement learning. Various applications of neural networks will be discussed.

The course will also cover deep learning libraries (e.g., Chainer, Tensorflow) and how to train neural networks using GPUs and GPU clusters. The course has a 20,000 hour allocation on Blue Waters which will be used for course projects.

Grading:

There will be homeworks (30%), a midterm (30%), and a project (40%).

Late Policy for homeworks:

0-2 days late: -10%

2-4 days late: -25%

More than 4 days late: homework will receive a 0 (i.e., no late submissions accepted after 4 days)

Office hours:

Biplab Deka: 5:30-7pm, Wednesday, Siebel Center 0207

Logan Courtney: 1:00-2:30 Tuesday, 21 Transportation (basement)

Justin Sirignano: 9:00-10:30 Wednesday, Transportation 209C

References:

"Deep Learning" by Goodfellow, Bengio, and Courville <http://www.deeplearningbook.org/>

Technical papers will also be provided as reading.

For basic background on machine learning: "The Elements of Statistical Learning" by Hastie, Tibshirani, and Friedman.