STAT GR5242, Advanced Machine Learning

Description:

Advanced Machine Learning is the third course in the Machine Learning sequence, culminating the skills and knowledge from Statistical Computing & Introduction to Data Science (GR5206) and Statistical Machine Learning (GR5241). The course will go into depth on deep learning and reinforcement learning, and will include assignments in the Python language. This course is made up of two parts: the first covers a wide range of deep-learning approaches to common ML problems, and aims to equip students with an advanced toolkit for real-world ML applications. The second part covers principles of reinforcement learning and deep learning approaches to reinforcement learning.

This course is accompanied by a companion course: GR5245 which covers Python and TensorFlow. It is recommended that all students take this course unless they have substantial experience with Python and TensorFlow.

Administrative details:

- Section 001: TR 10:10am-11:25am Location: 402 Chandler
- Section 002: TR 11:40am-12:55pm Location: 301 Uris Hall
- Instructors: John Cunningham and Kamiar Rahnama Rad
- TAs: Luhuan Wu, Charlie Windolf, and Jaesung Son
- Office Hours: to be announced
- Note: all instructors and TAs will teach both sections; there is no difference between the sections.

Evaluation:

- The course will involve six substantial problem sets, including programming work, mathematical problems, and data analyses.
- Homework submission will be detailed in a future lecture
- No late work will be accepted under any circumstances.
- To accommodate unexpected circumstances, we have implemented two important features:
    a. Your lowest assignment grade will be automatically dropped at the end of the term.
    b. You may submit and resubmit your assignment as many times as you like up until the deadline. This means that you should submit any partial solutions as you complete them, to make sure you receive as much credit as possible for the work you have done. After the deadline, the system will not allow you to submit anything. If you do not submit anything by the deadline, you will get a 0. There will be no exceptions to this rule. Submit your work early.

Topics:

The course will be taught in two halves, with the following tentative topics:

- Part I: deep learning
    a. Neural network basics and backpropagation
    b. Convolutional neural networks and residual networks
    c. Automatic differentiation
    d. Transfer learning
    e. Stochastic optimization
    f. Recurrent neural networks, attention/transformer networks
    g. Other topics: (as time allows) generalization / double descent; deep generative models (GAN, VAE); massive scale models (GPT3, DALL-E, OPT175); deep ensembles; embeddings

- Part II: reinforcement learning
    a. Introduction to RL and multi armed bandit
    b. Markov Decision Processes and Bellman equations
    c. Dynamic Programming and policy iterations
    d. Temporal difference learning and Deep Q Networks
    e. Double DQN, Prioterized experience replay and Dueling DQN
    f. Policy gradient methods
    g. Other topics: RL in psychology and neuroscience (1 lecture)

Prerequisites:

The course will be challenging to complete without first having a working knowledge of the following topics:

- Enough programming knowledge, including in Python, to use Tensorflow
- Basic unsupervised Learning, including k-means and PCA and the concept of a latent representation
- Basic optimization, including gradient descent and Newton's method.
- Supervised Learning, including multiclass logistic regression (and dependent topics) and support vector machines
- Exposure to models that address dependent structure in the data, such as time series models, Gaussian Processes, or HMM/Markov Chains