Artificial and Biological Neural Networks: Prof. Han Amarasingham: Fall, 2019, GC, Friday, 1:00-4:00 PM [CS 87100/PSYC 80103/BIO 79303]

**Course Description**

Artificial neural networks (ANNs) are pattern-recognition devices inspired in part by some simple facts we know about the brain. They now play a prominent role in artificial intelligence. But how "neural" are ANNs? How will developments in neurobiology bear on current and future developments of ANNs? We will survey modern applications of machine learning relevant to this question, and learn how to implement and execute the basic ideas. And we will review the attempts to connect, in a broader sense, contemporary ideas from machine learning to neurobiology. Such connections are currently a fertile area of research and debate. How are such questions informed by what we know about the brain experimentally, at the level of systems? Statistical and probabilistic ideas, including causal inference, are important here, and we will examine the conceptual issues these raise. We will read both classic and state-of-the-art papers on these topics, focusing on the core conceptual issues. The final project will either be a survey of some aspect of the literature on this topic, or a project oriented towards application of the ideas discussed in the course.

**Coverage Plan**

The course will be a combination of lectures and seminar-style discussion of research papers. In lectures, we will provide an overview of systems neuroscience (experimental techniques, interpretation of neurophysiology data, and overviews of theoretical neuroscience), an overview of major techniques in ANN's and their computational foundations (multilayer feedforward and recurrent networks, basics of statistical learning theory, backpropagation, convolution nets, implementation packages, and selected miscellanae from machine learning). Homework will focus on tutorial-style implementation (e.g., applications to image and language processing) and developing mathematical understanding. Lectures will be balanced by seminar-style reading in which students will be asked to lead and actively participate in discussions of relevant papers, which will range across statistics, neuroscience, computer science and philosophy disciplinary content. These readings will drill deeper into applications, questions of theory, and the biology of neural nets. Final projects will be developed over the course of the semester, most likely in groups, and presentations will wrap up the course.

**Prerequisites**

I would like to make the class accessible to doctoral students in the many fields such research questions touch on. (Groupwork therefore can be a mechanism for sharing expertise.) An excellent background for the course would be strengths in two of the following three areas, or some variation thereof: i) programming, ii) an undergraduate probability class, at the level typically taught in a math department (e.g., Math 375 in the case of CCNY), and some exposure

to statistics, and, iii) familiarity with systems neuroscience. *If you have questions about the above, please contact the instructor; we'd like to be accommodating.*