Practical Neural Networks
Instructor: Andrew Stanek, Junior in Biology

Course Syllabus

This is a course in practical Neural Networks. Though advanced Caltech CS courses provide a theoretical and highly rigorous treatment of machine learning, Practical Neural Networks would be a brief and instructive overview designed for casual programmers less interested in the nuances of the theory who simply want to include neural networks in their programming work. We will go over single and multi-layer, fully-connected perceptrons including simplified error analysis and backpropagation, starting at the most fundamental level. The course will also teach students about industrial and academic applications of neural networks and how to recognize tasks that neural networks are suitable for. Only a rudimentary knowledge of programming is required for this course.

The course text will be the online version Artificial Intelligence: A Modern Approach (Second Edition) by Stuart Russel and Peter Norvig with select excerpts as recommended reading (most of the text is beyond the scope of the course). To pass the class, students will need to write a multi-layer perceptron in the language of their choice. Each student should write a neural network for a specific purpose that he or she will choose. At the end of the class, the students will explain what their neural network is designed to do, and how (if at all) it varies from a standard perceptron. A single, extremely brief paper on Backpropagation will also be assigned as homework during week 7. The completed perceptron will be due at the end of finals' week. There will be two one-hour classes with an additional hour of homework a week. The first half of each class will be lecture, the second half will be given to group discussion and work on the neural network. To pass the class, students would be required to submit the paper on Backpropagation and complete the aforementioned neural network and demonstrate its functionality.

Topics by Week:

1. Introduction to the Course, History of Neural Networks, the Single Layer Perceptron
2. Single Layer Perceptron (*cont.)* vs. Multi-Layer Perceptron, Limitations of the System
3. How to Input to a Neural Network; Building the Training Set
4. Weights, Biases, Nodes, and Layers -- Putting Together the Single Layer Perceptron
5. Expanding to a Multi-Layer Perceptron
6. Error Analysis
7. Error and Training the Network; Backpropagation
8. Backpropagation (Cont.)
9. Optimization -- Non-Fully Connected Networks and Tricks for Speed
10. Wrapping Up -- Finishing the Neural Networks and Presentations.