ECE 3xxx: Neural Foundations of Machine Learning

GT Metz, Summer 2023

Objectives:

- Students will be able to demonstrate and build small unsupervised and supervised Neural Network algorithms on digital computers.

- Students will become familiar with the history of Machine Learning, the core single layer and two-layer supervised and unsupervised algorithms in machine learning, and with the neurobiological and physical computation foundations of Machine Learning.

This course provides a foundation for machine learning concepts, biological foundations, and implementation for using machine learning concepts as well as empowering students taking next level (senior and graduate level DSP or CS centric) machine learning courses.

Corequisite: Differential Equations (e.g. Math 2552); *Prerequisite*: Linear Algebra (e.g. Math 1554). Taking or having taken Physics 2 (Phys 2212) is encouraged, although not required.

Grading: Projects: 50%, Exams 30%, Final Project: 20%
- 4 Projects (one project is a short report)
- 2 Exams
- 1 Final Project (in lieu of a final exam)

3 Projects and the Final Project will involve computational experiments on your digital computer in MATLAB or Scilab or Python. Related program languages are possible with prior approval of the professor. The students will get to write their own Neural Network algorithms and demonstrate the properties of those algorithms. Projects will be a submitted .pdf file as a writeup of the project questions as well as include requested code.

The 4^{th} Project is a short report exploring a topic of the foundation in Machine Learning / Neural Networks.

The course has no required textbook. Materials will be made available on-line to the students.

This course will utilize one variant of an inverted classroom. Short lecture nuggets will be available and will be expected to be watched before class. Class time will be focused on interactive discussions, Q&A, and problem solving. Images of all whiteboards will be made available on-line. To not discourage interactive discussions, no recording of any form is allowed in the classroom for any reason.

Attendance in class is strongly encouraged, and class discussions may contain useful technical or administrative information. Students are also encouraged to read the GT catalog on attendance: https://catalog.gatech.edu/rules/4/.

Each student should govern themselves at least at the ethics of the Georgia Tech Honor Code (https://policylibrary.gatech.edu/student-life/academic-honor-code). Cheating will not be tolerated in this course, and will be referred to the dean of students office.

If a student has a disability to be accommodated for this class approved by the office of disability services (https://disabilityservices.gatech.edu), all reasonable attempts within the necessary constraints of the class will be made to find a solution. The student must identify these items the first week of class and contact the professor during the first week of class to discuss what ways a disability will be accommodated.

Course Outcomes:

After successfully completing this course, students will be able to
- Recognize core components and structures of neuro-inspired systems
- Analyze Neural Network systems for a particular application
- Design small Neural Networks using supervised and unsupervised learning

Course Outline:

<u>Fundamentals for Machine Learning</u>
Linear ODEs
Basic Linear Circuits
Other preliminaries

<u>Neuroscience for computation</u>
Biological membranes and Channels
Neurons
Synapses and learning
Cables and dendrites
Layer of Neurons, Winner-Take-All Networks
Layers of Neurons (e.g. Huber and Wiesel)
Human Cortex and Neural Structures

<u>History of Machine Learning & concept introduction</u>
Pre 1980s (Perceptrons, Adaptive Filters)
First Neural Network wave (Hopfield networks, Energy Surfaces, Backpropagation)
First Killer applications & established techniques
SYNAPSE and Deep Learning wave

<u>Unsupervised Learning - computation on Statistics</u>
Matrix of linear ODEs
Required Statistics: mean, covariance, erogodicity
SOM (self organizing map), VQ, clustering
Unsupervised layers & Map formation
Oja rule and normalized solutions (e.g. PCA, ICA)

Supervised Learning

- Adaptive filters, LMS, Perceptrons
- Backpropagation
- Deep networks & many layer networks

Issues in learning

- Universal approximation (function approximation)
- Classifier theory (e.g. ROC curves) and Hyperplanes
- Generalization and overfitting

Introduction to Physical Implementations

- Parallel computation introduction, computing on mesh
- Computing in Memory, Nonvolatile memory, and further topics