ECE 3xxx: Neural Foundations of Machine Learning

GT Metz, Fall 2025

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.

Prerequisite w/ concurrency: 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.

Sponsoring Technical Interest Group (TIG): BioEngineering

Grading: Projects: 70%, Exams 30%,
- 5 Projects. The last project takes the place of the final exam, and that last project will be due during final exam period.
- 2 Exams (week 4, 10)

The five Projects 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.

Project Writeup: Projects will be written up in double-column, single spaced, IEEE format. One can find alot about this format on-line (e.g. IEEE), and it is the one format you need to learn as an ECE individual. Headings given in typical IEEE format templates are suggestions only; you will need to have at least a few sentances explaining the project, and you will need to have discussion throughout your project. Figures must be included in the middle of the writeup. Discussion of your results is essential and always expected even if not explicitly stated.

Graduate students taking this course for Graduate Credit: For Graduate students wanting to take the course for Graduate Credit, two additional project assignments are required that will be related to their graduate focus
- One project related to unsupervised learning by November 6th
- One project related to supervised learning by December 6th.

The resulting grading in that case will be:

Projects: 75%, Exams 25%,

- 7 Projects (last project due during final exam)
- 2 Exams (week 4, 10)

Graduate students taking this course will have a separate grading curve from the undergraduate students where that might be required.

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

Course Structure: 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: Attendance in class is required, 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/.

Collaboration and Cheating: You may discuss the questions among students in the class, but each person must independently perform and write-up the required work.

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 any case of cheating will be referred to the dean of students office.

The professor reserves the right to use software to check for plagiarism of material. Further, as the assignments are different than typical Machine Learning class assignments, one should be careful about taking material from other websites, particularly without giving proper reference to that source. If you attempt to use an AI code assistant, you must acknowledge that use as well, although realize the code assistant might give you incorrect code, which is entirely the student's responsibility.

Disability accomodations: 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 will be made to find a solution if it is possible. 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 ML systems
- Analyze Neural Network systems for a particular application
- Design small Neural Networks using supervised and unsupervised learning
- Understand basic neuroscience that provides the foundation for Machine Learning approaches.

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 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)
- Saliency and attention

<u>Supervised Learning</u>
- Adaptive filters, LMS, Perceptrons
- Backpropagation
- Convolutional networks
- Deep networks
- Introduction to LLM networks

<u>Issues in learning</u>
- Universal approximation (function approximation)
- Classifier theory (e.g. ROC curves) and Hyperplanes
- Generalization and overfitting

<u>Introduction to Physical Implementations</u>
- Parallel computation introduction, computing on mesh
- Computing in Memory, Nonvolatile memory, and further topics