An AI Course in Computational Engineering at UT Austin

Joe Stubbs, PhD
Sept, 2025

About the Speaker

- Lead, Cloud and Interactive Computing Group, TACC, UT Austin
  - 15 full-time staff plus students
- Faculty appointment, Computational Engineering
- Co-PI of NSF-funded ICICLE AI Institute
- Interests: Cloud computing and distributed systems, formal methods and machine learning

COE 379L: Software Design For Responsible Intelligent Systems

Course Overview

- Objective: Introduce students to scalable data analysis and machine learning.
- Focus on techniques for designing, implementing, validating and operating responsible intelligent systems – how do they know if their system will be reliable?
- Projects-based course with “real” software development at the core.
- All materials available online, e.g., https://coe379l-sp25.readthedocs.io/en/latest/

History

- Initially taught in Spring of 2024 with about 20 students
- In Spring 2025 we gave it a second time and it had 30 students plus ~110 on the waitlist.
- Currently teaching the third iteration (Fall 2025) with about 35 students.

Motivation for the Course

1. COE 332 (Software Design for Cloud and Distributed Systems) very popular, and many students often requested “more”.
2. With COE 332, the goal was partially self-fulfilling: what skills did students need to be successful in an undergraduate research experience project (the kinds that my grants had funds for).
3. With the ICICLE AI Institute we were doing more AI/ML work and had funding for students to work on those projects.
4. We wanted an excuse to learn more about certain topics ourselves! :)

COE 379L: Overview and Topics

Material

- Unit 1: Data Analysis and Data Preparation
- Unit 2: The Machine Learning Application Development Life Cycle and Classic Machine Learning Algorithms
- Unit 3: Neural Networks and Deep Learning
- Unit 4: Advanced Topics (Transformers, Benchmarks, RAG, TinyML, etc)

Software

- Unit 1: Numpy, Pandas, Matplotlib, Seaborn
- Unit 2: Scikit-learn
- Unit 3: Tensorflow Keras, Docker and HTTP servers (e.g., Flask)
- Unit 4: Transformers, LangChain

Projects

- Project 1: Breast Cancer classification (solo)
- Project 2: Housing Prices (solo)
- Project 3: Predicting Damage to Buildings in Images pairs)
- Project 4: Open-ended (pairs)

Interactive, Hands-on Lectures

- Every student has their own VM running in our cloud with access to required software
- Lectures consist of a series of presentations followed by a hands-on portion that reinforces the concept
- Students are given time throughout the lecture to work on these hands-on portions
- We aim to give students the opportunity to explore the topics with open-ended questions and gain confidence in using the tools

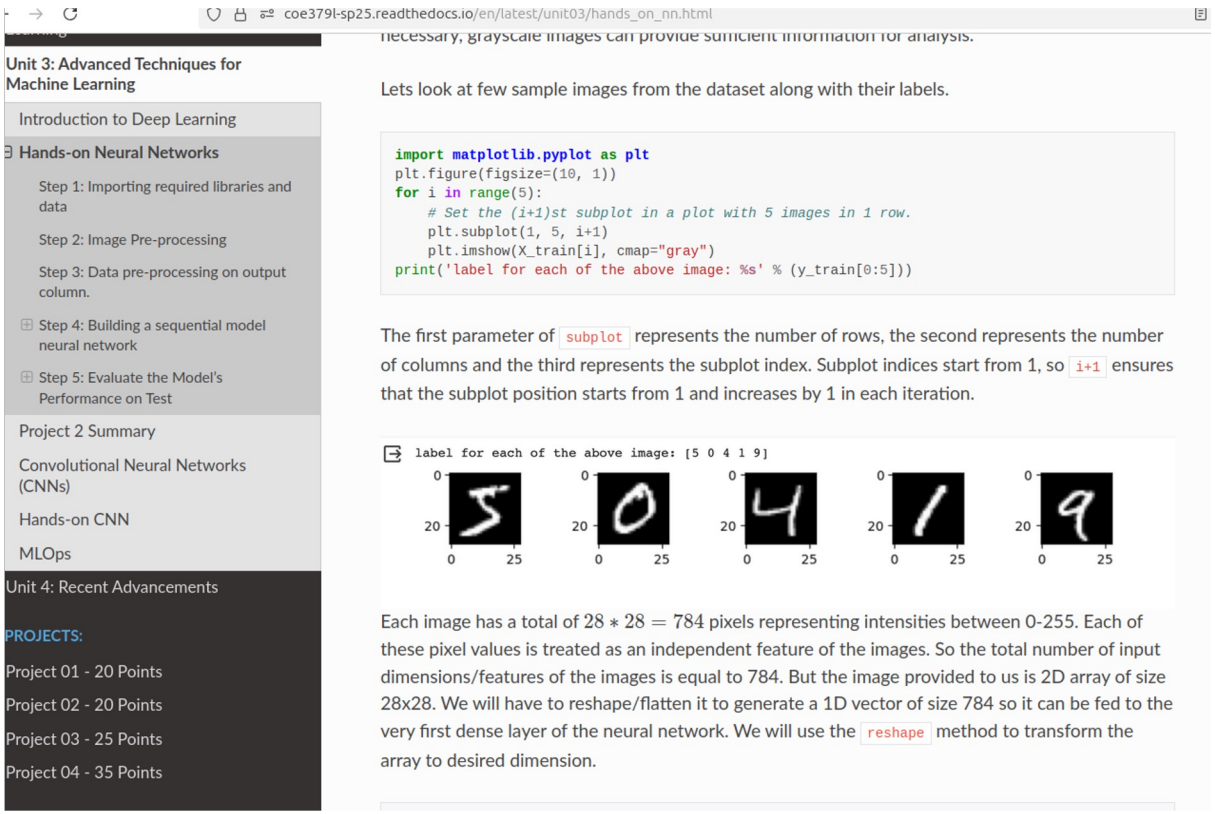

Project Assignments

- Analyze a real-world dataset using one or more machine learning methods.
- Do the necessary pre-processing, model design, training, and validation.
- In some cases, we have them read papers from the literature and ask them to incorporate ideas from the paper into their work.
- In some cases, we ask them to develop deployment artifacts (e.g., containers) so that code can be used
- Write a report describing the analysis they did, the choices they made, the results they obtained, etc.

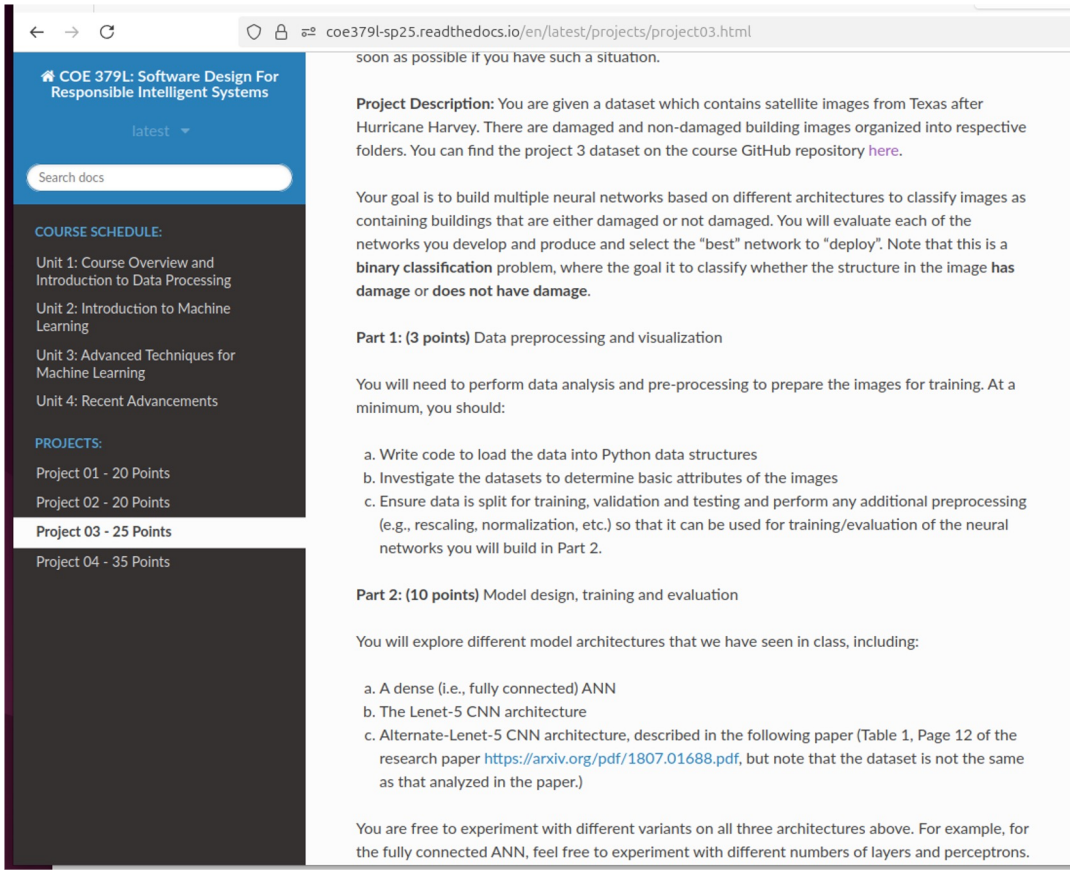

Use of AI

Goal

- Want students to learn to use AI tools **and** to learn the actual AI/ML topics.
- AI can be a great tool for *learning* when used correctly (e.g., AI tutors)
- AI can also be a great tool for *doing* development – but users must be able to recognize when the AI gives incorrect, misleading or otherwise suboptimal solutions.

In Practice

- Students will use the AI to solve an entire project – it's just too tempting.
- Students don't understand when the AI makes a "mistake" (or when the student tries to use an AI-generated solution out of context)
- Some students use of AI prevents them from learning the material.

AI Policy and Written Exam

Use of AI Policy:

- Every use of AI that leads to new code must be documented in a “Use of AI” file – like a bibliography that contains the tool, prompt used, and code generated.
- Within the code, must add (in a comment) a reference to the Use of AI document anywhere that code was generated.
- Do not use AI to generate the reports.

Exam:

- This semester we will have a written exam covering conceptual ideas from the course.
- The goal of the exam is to make students aware that we want them to learn the material, not just have AI do everything.