Special Topics: Machine Learning Applications in Experimental BME Research

(42-698)

Spring 2026

Lecture: Monday and Wednesday, 12:30–1:50 PM
Location: GHC 4211

Instructor and TA

Instructor	Prof. Newell R. Washburn
Email	washburn@andrew.cmu.edu
Office Hours	Monday, 2:30–3:30 PM, Scott 4N203
Teaching Assistant	Tushar Nayak
Email	tusharn@andrew.cmu.edu
Office Hours	Wednesday, 11:00 AM–12:00 PM; Friday, 1:00–2:00 PM

Zoom Information

When needed for lectures or office hours, the course will use the following Zoom meeting link:

Join Zoom Meeting

Meeting ID	999 1249 4785
Passcode	433379

Course Description

This course introduces students to applications of machine learning and artificial intelligence in experimental biomedical engineering research. The course is organized into four major sections:

(1) Traditional approaches to machine learning, including featurization, regression, and classification;

(2) Neural networks and deep learning, with applications to time-series and visual data;

(3) Transformers and attention mechanisms;

(4) Large language models and retrieval-augmented generation.

The course focuses on common data types generated in experimental biomedical engineering research, including tabular data from sets of experiments, image data, spectral data, and time-series data. A diversity of regression and classification methods will be introduced, including linear models, Gaussian processes, tree-based methods, and TabPFN. Particular emphasis will be placed on experimental design and modeling datasets with small sample sizes.

Important examples will include quantitative analysis of cells in culture, in situ spectroscopic data, and omics data. Methods of statistical analysis, feature selection, and transfer learning will also be introduced. Students will complete a final project that allows them to apply these tools to a topic of their choosing, preferably related to their own research.

In parallel with standard machine learning tools for regression and classification, this course will introduce recent advances in large language models and AI agents. In the class project, students will have the opportunity to apply retrieval-augmented generation and tune large language models that can support the specific modeling methods developed in the course.

Prerequisites

Graduate standing or consent of the instructor is required. Familiarity with Python programming is expected.

Course Outline

Part 1: Machine Learning Algorithms for Classification and Regression

- Random forests, XGBoost, and related tree-based methods
- Gaussian processes
- TabPFN
- Problem Set 1

Part 2: Neural Networks and Deep Learning

- Artificial neural networks
- Recurrent neural networks and LSTMs with applications to time-series data
- Convolutional neural networks with applications to visual data
- Problem Set 2

Midterm Examination

Part 3: Transformers

- Attention mechanisms
- Transformer architectures
- Applications of transformers
- Problem Set 3

Part 4: Large Language Models

- Large language models
- Retrieval-augmented generation
- AI agents and related tools

Grading

Final grades will be determined according to the following breakdown:

Component	Weight
Problem Sets	33%
Midterm Examination	33%
Final Project	34%

Final Project

Students may choose their final project topic with instructor approval and are encouraged to base the project on their own research. The final project will require the following components:

- A written report;
- An oral presentation in class;
- Submission of source code.

All components must be completed. Failure to complete any required component will result in a score of zero for the final project. Oral presentations will be conducted during the scheduled final examination period, and attendance is mandatory.

Policy on Use of AI

Students are encouraged to use online and AI-based resources to assist with coding the algorithms needed for the problem sets. However, all written answers to problem-set questions and all final-report text must be written by the student.

Written submissions may be checked using common AI-detection tools. Students are encouraged to check their own work prior to submission. Positive identification of AI-produced text, defined as greater than 50% likelihood, will result in a score of zero for that submission.

Textbook and Related Materials

There is no required textbook for this course. Resources and related materials will be shared by the instructor.

Academic Integrity Policy

Students are expected to follow Carnegie Mellon University's academic integrity policies. The policy is available at:

http://www.studentaffairs.cmu.edu/acad_int/

Important Dates

The following dates are tentative and subject to change:

Date	Item
February 11	Problem Set 1 due: Machine Learning
February 27	Problem Set 2 due: Deep Learning
March 2–6	Spring Break
March 18	Midterm examination, in class
April 3	Problem Set 3 due: Transformers
April 20–22	Class presentations
May 4	Final projects due for graduating students
May 9	Final projects due for non-graduating students

Learning Outcomes

By the end of this course, students should be able to:

1. Understand the principles of machine learning.
2. Apply machine learning and AI tools to model experimental data.
3. Model experimental data using a diversity of machine-learning tools.
4. Implement new machine-learning techniques through self-study.
5. Understand neural network architectures and apply deep learning methods.
6. Understand transformer architectures and their applications to modeling.
7. Understand how to tune large language models using retrieval-augmented generation.