

Modern Large Language Models (STAT 4720/5720)

Spring 2026

INSTRUCTOR

- Dr. Gavriil Shchedrin, Lecturer in APPM with over 4 years of industry experience as a Lead Data Scientist.
- LinkedIn: <https://www.linkedin.com/in/gavriils>
- Email: gavriil.shchedrin@colorado.edu

COURSE GOALS

- This course provides a hands-on introduction to modern Large Language Models (LLMs) using PyTorch and TensorFlow. Students will gain practical experience with the technical stack that is essential for developing and deploying large-scale machine learning models, equipping them with all the necessary skills required in the modern DS/ML.

The course is structured around two major projects: building a LLM from scratch and fine-tuning an existing LLM using pre-trained models from open sources e.g. Hugging Face. Through these projects, students will explore the key deep learning concepts, including but not limited to transformer architectures, attention mechanisms, tokenization, optimization techniques, and soft fine-tuning.

Students will also develop expertise in model training, performance evaluation, and hyperparameter tuning. The fine-tuning project will introduce transfer learning techniques, allowing students to adapt powerful pre-trained models for specific applications while optimizing efficiency and performance of the respective models.

The class is structured based on the cooperative learning model known as pair programming. In this learning model, students will be divided into driver-navigator pairs where the driver actively codes and implements solutions, while the navigator provides guidance, reviews the code, and helps with debugging the code. This approach fosters teamwork, enhances problem-solving skills, and ensures that students are constantly engaged with both theoretical and practical aspects of the material.

By integrating industry-relevant tools with collaborative learning techniques, this course will prepare students to tackle real-world challenges in deep learning and natural language processing, setting students for a successful career in modern DS/ML.

COURSE PREREQUISITES

- Math: Multi-variable Calculus, Linear Algebra (solid intermediate level, equivalent to APPM 3310), knowledge of regression methods and statistics.
- Coding: Python (solid intermediate level)
- This will be a fast-paced course geared towards graduate students and advanced undergraduates with a solid background in math and coding.

COURSE STRUCTURE

- There will be weekly/bi-weekly homeworks in the form of coding problems. Additionally there will be one in-class midterm exam and an in-class final exam in a form of a real-world technical interview, administered by an instructor.

HOMEWORKS

- The homeworks will be based on a small-to-medium size problems solving of which will give a student a mastery of the data science tools.
- **Disclaimer:** using publicly available LLM models (e.g. chatGpt, Claude, etc) is allowed by the students with the goal of *refining* their solutions to the coding problems. However, if students merely copy-paste LLMs solutions they will totally miss the goal of the class, which is to equip students with the modern data tools and core technical skills that are required to be successful in the data science field.