Theoretical Foundations in Modern Machine Learning

CSCI-GA.3033-141 | Fall 2026

Instructor	Allen Liu	**Class time**	Tuesdays, 4:55–6:55 PM
Location	Bobst Library, LL138	**Course website**	aliu42.github.io/fall26-topicsml.html

1 Course description

This class aims to showcase various ways that theoretical insights play a role in advancing modern machine learning. We will focus on theoretical perspectives and simple mathematical models that both help us understand modern training pipelines and inform new algorithms and design principles. Topics include scaling laws, data and compute allocation, post-training and reinforcement learning, and inference-time interventions.

Organizing principle

For each topic, we will use theory to move from a practically important question to tractable mathematical models that help us make concrete predictions and design decisions. We will see that despite the empirical nature of modern AI research, these mathematical principles can provide surprisingly strong guidance.

2 Prerequisites

Students should be comfortable with probability, linear algebra, and multivariable calculus. It will also be helpful, although not required, to have some familiarity with deep learning and neural networks.

3 Course policies

The course grade will be based on class participation and two to three in-class presentations spread throughout the semester. The final number of presentations will be determined after enrollment stabilizes.

Component	**Weight**
Class participation	20%
Two to three in-class presentations	80%

Each presentation will focus on understanding a research paper or a research direction represented by a few closely related papers, and on presenting the key ideas to the class. Presentations will be roughly 20 minutes, including questions and discussion. Participation grade will be based on active engagement during lecture and with other students' presentations.

AI usage

The use of AI tools is permitted and encouraged. Apart from the primary content of this class, I think it is also extremely important to learn how to work effectively with modern AI tools.

AI tools can be very helpful for understanding technical papers and preparing presentations. Students are nevertheless responsible for the accuracy of what they present and for understanding the material themselves. The ability to explain the key ideas and answer questions during a presentation will provide a direct check of that understanding. A great quote: ***"You can outsource your thinking, but you cannot outsource your understanding."***

4 Course organization

LLM fundamentals (Meetings 1–2)

Autoregressive modeling, training objectives, transformer fundamentals, and an overview of the modern language-model training pipeline.

Pre-training (Meetings 3–5)

Scaling laws, data and compute allocation, data quality and mixtures, and the systems constraints that shape scaling decisions.

Post-training (Meetings 6–9)

Reinforcement learning, policy gradients, group-relative methods such as GRPO, reward and value modeling, and preference learning and optimization.

Test-time scaling (Meetings 10–12)

Parallel sampling, best-of-N selection, search methods, self-verification, and adaptive allocation of inference-time computation.

Flex topics (Meetings 13–14)

The final topics can be selected potentially based on the interests of the class. Possibilities include alignment, interpretability, distributed training, diffusion models, or other recent developments related to the core themes.