LLMs for Computational Linguistics

Course Number: LING 4431
Lecture: Monday and Wednesday, 3:30 - 4:45
Location: White Gravenor, 209

Instructor: Ethan Wilcox
Office: Poulton 248 (Office Hours: Wednesday, 2-3 pm)
Email: ethan.wilcox@georgetown.edu
Website: https://wilcoxeg.github.io/

TA: Ping Hei Yeung
Office Hours: TBD
Email: py101@georgetown.edu

Summary: Large language models (LLMs) are the foundational technology behind today's most advanced artificial intelligence systems. They have revolutionized the field of natural language processing and challenged traditional ideas about how language is represented in the human mind. This course offers an introduction to LLMs from the perspective of computational linguistics. Through lectures and hands-on demonstrations, we will explore three interrelated questions: How do LLMs work at a technical level? How can they be used to process natural language data? And how can they be used to model human linguistic cognition? The first half of the course will cover technical foundations of LLMs, including the transformer architecture, tokenization, assessment techniques, and scaling. The second half of the course will focus on LLMs' applications in natural language processing, linguistics, and cognitive science. Topics may include security and privacy, ethical issues, multilingualism, and LMs as models of human language processing and acquisition. This course is appropriate for advanced undergraduates and graduate students. Students will gain experience training (small) language models, implementing basic interpretability techniques, and reading recent research papers in the area. While there are no formal prerequisites, knowledge of programming in Python and basic math and statistics will be assumed.

Learning Outcomes:

- Students will learn the technical foundations of LLMs and gain experience training (small) models.
- Students will learn the techniques for analysis and evaluation of models, and gain experience running evaluation methods on small models.
- Students will gain experience thinking critically about the technical limitations of LLMs, as well as their ethical drawbacks, including risks to privacy and potential harms.
- Students will learn how LLMs are changing the fields of linguistics and cognitive science, and gain experience reading current research papers on the use of LLMs as cognitive models.

Course Structure:

- This course is a lecture course, and will meet twice a week for 1h 15m
- This course will include a pen-and-paper final exam

Course Schedule

Week	Module	Topic	Assignments
Jan 5, 2026 (1 session, Jan 7th)	Introduction	* Course overview & Introduction	
Jan 12, 2026	Technical Foundations	* What is an LLM? * How to evaluate an LLM: Benchmarking	1/12 Annotated Reading 1 out
Jan 19, 2026 (1 session, Jan 21st)		* Benchmarking & LLM Capeabilities	1/19 Annotated Reading 1 due 1/19 Coding Assignment 1 out
Jan 26, 2026		* Technical Foundations: Linear algebra & deep learning * Transformer architecture	
Feb 2, 2026		* Transformer architecture	2/2 Coding Assignment 1 due 2/4 Coding Assignment 2 out
Feb 9, 2026		* Training Transformers: Gradient Descent	
Feb 16, 2026 (Monday session held on Tuesday, Feb 17th)		* Post-training: Instruction tuning and RLHF * Tokenization	
Feb 23, 2026		* Scaling and Parallelization	2/25 Coding Assignment 2 due
Mar 2, 2026		🌴 Spring Break 🌴	
Mar 9, 2026	Topics in NLP	* Prompting * Generation	3/9 Coding Assignment 3 out
Mar 16, 2026		* Information retrieval * Multilingualism and low-resource NLP	
Mar 23, 2026		* Ethics: Security and privacy * Ethics: Biases and harms	2/23 Coding Assignment 3 due 3/23 Annotated Reading 2 out
Mar 30, 2026		* Ethics: Environmental impacts of LLMs * Group Discussion: Ethics	4/1 Annotated Reading 2 due
Apr 6, 2026 (1 session, April 8)	Linguistics & Cognitive Science	* LLMs as models of human language processing	
Apr 13, 2026		* 4/13: Field trip: Law and corpus linguistics book talk (4:00 pm) * LLMs as models of human language acquisition	4/13 Annotated Reading 3 out

Apr 20, 2026		* Speech LLMs and Linguistic Theory * Group Discussion: Implications for linguistics and cognitive science	4/22 Annotated Reading 3 due
Apr 27, 2026 (1 session; April 27th)		* Recap & final exam review	
May 2, 2026			**Final Exam, 4:00 pm**

Assignments: Throughout the semester, you will have six assignments – three coding assignments and three annotated readings

- **Coding Assignment 1 (Benchmarking):** You will be provided with code that evaluates an LLM on a benchmark. Your job is to evaluate the model on the benchmark with different prompts. How does the choice of prompt change the model's performance?
  What you need to turn in: Writeup justifying prompt choice; model results on benchmark

- **Coding Assignment 2 (Training a Transformer from Scratch):** You will be provided with code that implements a transformer inside of a Google Colab notebook. Your job is to train the model and run code to evaluate it on a benchmark.
  What you need to turn in: Writeup about the training process; results of evaluation

- **Coding Assignment 3 (Dataset Hacking):** Students will repeat the process in assignment 1. However, before re-training the model, they will make one change to the dataset. Options for this change include: (a) re-ording the dataset based on some criteria (i.e., create a data curriculum); (b) remove some feature of the dataset (e.g., sentences longer than a certain length); (c) some other change of your choice. Students will repeat the same evaluation and report the results as part of a short write-up.
  What you need to turn in: Writeup about the change and rationale, results of evaluation

- **Annotated Readings:** In addition to the assignment, students will be required to turn in 3 annotated readings. These can either be PDFs or printed papers; for the last two I'll also ask for questions that we'll use during our group discussion. Annotations can be of various formats, but must demonstrate active engagement with the text. Annotations can include: highlighting of key passages, in-line questions, in-line comments. Readings will include:

  - Reading 1: Capabilities of LLMs
  - Reading 2: Ethics of LLMs
  - Reading 3: LLMs, Linguistics and Cognitive Modeling

All submissions happen through Canvas.

Late Policy: Late homework can receive credit, but the grade you can achieve is capped based on how late it is, with the maximum grade you can achieve lowered by 5% per 24-hour period after the due time for the first two days, and 10% per 24-hour period thereafter. You can always receive up to 40% of the full grade if you turn in an assignment before the final day of class.

Group Work Policy: The text that you submit as part of homework assignments *must be written by you*. However, working in groups on assignments can be extremely helpful. If you work with a group to come up with an overall plan or *pseudocode*, or if you work with a group to understand and debug a piece of non-functioning code, that's OK. But you need to implement the pseudocode, or fix the bug yourself!

Final Exam: This class will include a pen-and-paper final exam. The focus will not be on specific formulae, but rather on your ability to solve problems, answer questions about high-level concepts covered in class, and demonstrate critical thinking about LLMs and their applications.

Participation and Attendance: This course is a fast-paced introduction to a very complex topic. Regular attendance in class will ensure that you don't fall behind. Class sessions should only be missed for things such as planned medical events, medical emergencies, family emergencies, or religious observations. While we will not be keeping track of attendance formally, 20% of your grade will be assigned based on your participation, with 10% specifically allotted for participation and preparation for the two scheduled group discussions. Participation can take many forms – asking questions in class, engaging in group activities, and attending office hours. For more information regarding the university policy on attendance, please see the academic standards section of the Undergraduate Student Bulletin (https://bulletin.georgetown.edu/regulations/standards/)

Prerequisites:

- Students should have a basic ability to program in Python
- Students should be comfortable with basic statistics and algebra

AI Policy: Given that this class is *about* LLMs, I will not be limiting AI use and even encourage you to use it. However, as this class will make clear, AI is limited in many key ways, and I encourage you to use these tools with a critical mindset.

Grade Breakdown:

Final Exam	30%
Coding Assignments	25%
Reading Assignments	25%
Participation and Attendance	20%

Accommodations: If you have a recognized accommodation through ARC, please contact Ethan. More information about accommodations and support, including student-athlete support, can be found on the ARC website (https://academicsupport.georgetown.edu/)