Syllabus for CSCI 456-01, Large Language Models, Spring 2026

Instructor
Robert Michael Lewis
ISC 4, rm 1390
rmlewi@wm.edu
Office hours: M 0930–1130, T 0930–1130, or by appointment

Teaching Assistant
Wenhan Lyu
ISC 4, rm ????

The course In the words of the Grok 3 large language model (with one minor edit):

> This course dives into the electrifying realm of large language models (LLMs), the engines driving today's natural language processing revolution. Students will master the theory, architecture, and real-world applications of LLMs, preparing to harness–or perhaps politely welcome–our future AI overlords. Blending lectures, hands-on coding, and critical analysis, this course equips you with the skills to build, apply, and interrogate these linguistic titans.
>
> We kick off with a deep dive into the transformer architecture, peeling back the layers of self-attention, positional encodings, feed-forward networks, and multi-head attention. You'll see how these pieces scale into massive, billion-parameter LLMs, exploring pre-training, fine-tuning, and the gritty details of optimization (think gradient accumulation and mixed precision).
>
> Then, we shift gears to LLM applications, tackling hands-on topics like:
>
> - Prompt engineering: Shaping inputs to bend LLMs to your will—generation, classification, or reasoning.
> - Semantic search: Using embeddings to unearth answers from vast text troves.
> - Text summarization: Distilling documents with flair, abstractive or extractive.
> - Conversational AI: Crafting bots that chat like humans (or better).
> - Ethical deployment: Navigating bias, fairness, and the carbon footprint of our silicon masters.
>
> Using tools like Hugging Face's Transformers, students will experiment with LLMs through projects—maybe a custom search engine or a fine-tuned model for your favorite niche. Expect coding, paper discussions, and a capstone that ties it all together.
>
> We'll ponder the big stuff: How do LLMs reason? Where do they stumble? And are we ready to bow to their brilliance? By the end, you'll grok the tech, wield it like a pro, and maybe even tip your hat to our AI overlords. We don't just study the future, we salute it. Strap in, humans—it's time to decode the language matrix!

How to communicate with me The best way to communicate with me is to talk at me in person. Catch me after class or during my scheduled office hours or make an appointment. Please use email only for questions that require only a short answer (e.g., "yes", "no", "42"). If you do send me email, please let me know that you are in CSCI 456 so I will know which of my classes you are in.

CS system account You will need an account on the CS system for this class to take advantage of the GPUs on the CS machines. If you are a declared major any existing CS account should be retained but you should check to make sure. If you do not have a CS account you can request one at https://accounts.cs.wm.edu. CS accounts are not automatically created so expect to wait a couple of days.

Textbooks We will use several texts, all available online via Swem Library.

- For the transformer architecture, Build a Large Language Model (from Scratch) by Sebastian Raschka;
- For techniques for applying LLMs, Hands-on Large Language Models: Language understanding and generation, by Jay Alammar and Maarten Grootendorst;
- For natural language processing,Speech and Language Processing (3rd ed. draft), by Dan Jurafsky and James H. Martin;
- For some high-level topics, Large Language Models: A Deep Dive: Bridging Theory and Practice by Uday Kamath, Kevin Keenan, Garrett Somers, and Sarah Sorenson.

Version: January 21, 2026, 1518.

In addition I will post links to papers for us to read.

Digital etiquette Put away your cellphone, out of sight (or better yet, out of reach) once class begins. If you need to use your phone, or are a telephonically incontinent phone fondler, please step out of the room. If you do choose to fiddle with your phone in class, expect to be rusticated for the remainder of the class, if not longer.

On days where we do not have in-class activities that require laptops, please sit on the two sides of the classroom (not in the middle) if you are going to use your computer. This will reduce the distraction for others.

Laptops Laptops will be used frequently in class. Be sure to bring yours to class with a full charge (or a really long extension cord).

Where to find stuff

- Blackboard.
- Gradescope.

Important dates Academic calendar:

- Add/drop deadline: January 30.
- Midterm grading period: March 2–22 .
- Spring Break: March 7–15.
- Withdrawal deadline: March 23.
- Last day of classes: May 1.
- Final Exams: May 4–8, 11-12.

Rescheduling or deferring a final exam requires a petition on your part to either the Dean of Undergraduate Studies or the Dean of Students, depending on the situation:

> http://www.wm.edu/offices/registrar/calendarsandexams/examschedules/rescheduledeferexam/

Faculty do not have the authority to reschedule a final exam without authorization from the Dean of Undergraduate Studies or the Dean of Students.

Grading The final grade for this course will be based on the following.

- In-semester projects and assignments: 80%.
- Final project 20%.

Academic accommodations of disabilities Per Student Accessibility Services,

> William & Mary accommodates students with disabilities in accordance with federal laws and university policy. Any student who feels they may need an accommodation based on the impact of a learning, psychiatric, physical, or chronic health diagnosis should contact Student Accessibility Services staff at 757-221-2512 or at `sas@wm.edu` to determine if accommodations are warranted and to obtain an official letter of accommodation. For more information, please see http://www.wm.edu/sas.

Once I have received an accommodation letter, it is your responsibility to discuss accommodations and accommodation logistics with me.

Absences If you are absent for good cause (e.g., illness, a job interview, as a representative of the College in a musical tour) I will work with you concerning anything you may have missed. On the other hand, "not feeling it" is not an acceptable excuse for being absent.

The Honor Code Keep in mind all work in this course is subject to the College's Honor Code. Honor Code violations may result in severe penalties, including

- permanent dismissal from the College,
- suspension from the College,
- a failing grade,
- a grade reduction in the course, and/or
- having your nose torn off with red-hot pincers.

So if you have any questions about whether something might constitute an Honor Code violation, please ask me before you do it!

That said, I encourage you to work with others currently enrolled in CSCI 456 on the homework and projects. In addition, you may consult me, books, papers, or papers, including pre-prints. In the spirit of the course you may also confer with LLM/AI such as ChatGPT, Claude, Gemini, and Grok, and, unless otherwise noted, ask them for help, even to generate code for you.

You must state in your work what sources you have consulted, with whom you have collaborated, and from whom you have received help. Be sure to give any sources (including LLMs) proper credit in your write-up of your work.