

Course Outline: CS 690 007
Evaluating Generative AI Systems
Fall 2025, Tu 4:30pm, IN 132, Fairfax

Course Overview

An empirical look at systems that use generative artificial intelligence (GenAI). The first few weeks of the course will focus on methodology for evaluating GenAI, viewed through the lens of common use cases like task planning, reasoning, and computer vision. The second half of the course will focus on systems-level evaluation methodology, where individual system components interact with each other to create system-level behavior that is greater than the sum of its parts. Systems-level topics include healthcare, human-agent interaction, and robotics.

Prerequisite: Graduate student standing.

Instructor Name and Contact Information

David Porfirio, Ph.D

dporfiri@gmu.edu

Office: FUSE 7310

Office Hours: by appt.

Learning Outcomes

This course will provide the fundamentals of empirical evaluation of GenAI system performance. The primary vehicles for delivering these outcomes will be: (a) a course project, in which you will build a small system of your own interest, make claims about your system, and evaluate your system to support these claims; (b) weekly lectures on evaluation methodology; and (c) weekly readings and class discussions that focus on how this methodology is currently applied to cutting-edge GenAI research.

After taking this class, you should be able to:

- Understand how GenAI is applied in a variety of different applications and use cases
- Apply GenAI to these applications and use cases
- Understand different ways that the performance of individual GenAI models can be measured
- Understand the fundamentals of and how to apply experimental design to the evaluation of GenAI models and GenAI-infused systems
- Understand the principles of and how to apply statistical analysis to the assessment of GenAI model and GenAI-infused system performance
- Differentiate component-level vs. systems-level evaluations
- Understand the fundamentals of human-subjects research
- Apply novel GenAI technology to the creation of a system that humans can use
- Continue applying these learning outcomes even as GenAI so rapidly evolves

Course website

Course materials will be hosted on Canvas.

Code will be hosted here: <https://github.com/dporfirio/CS690-Fall25>.

Course Project

Students will work individually or in small groups on an applied-AI project. The project should: (a) intersect your individual interests with the course learning outcomes; (b) apply GenAI to a systems-level use case; and (c) apply your knowledge of evaluation methodology to the evaluation of your system.

Course Schedule

Readings and assignments will be posted on the course website.

8/26/2025	Introduction: GenAI Systems Course Overview
9/2/2025	Introduction: Methodology Overview Common GenAI Evaluation Methods
9/9/2025	Experimental Design I <i>Technology Focus</i> : Task Planning
9/16/2025	Experimental Design II <i>Technology Focus</i> : Reasoning Models
9/23/2025	Statistics I <i>Technology Focus</i> : VLMs
9/30/2025	Statistics II <i>Technology Focus</i> : Healthcare
10/7/2025	Systems Research <i>Technology Focus</i> : Robotics
10/14/2025	HSR I: Ethics <i>Technology Focus</i> : Human-AI Interaction I
10/21/2025	HSR II: Measuring Humans Quantitatively <i>Tech. Focus</i> : Human-AI Interaction II
10/28/2025	HSR III: Measuring Humans Qualitatively <i>Tech. Focus</i> : Human-AI Interaction III
11/4/2025	Combining Different Sources of Data <i>Technology Focus</i> : Social Robotics
11/11/2025	Longitudinal and In-The-Wild Studies <i>Technology Focus</i> : Home Robots
11/18/2025	Research Through Design: Looking to the future
11/25/2025	Project workday
12/2/2025	Final project presentations
12/9/2025	Project writeups due

Grading Policy

Grading Schema

A	92%	B+	87%	C+	77%	D	60%	F	<60%
A-	90%	B	83%	C	72%				
		B-	80%	C-	67%				

Grade Weights

In-class discussion participation - 12%

Project Milestones - 48%

Final Project Presentation - 15%

Presentation Peer Evaluation - 5%

Final Project Writeup - 20%

AI (Artificial Intelligence) Tools Policy

Generative AI is an important tool for students to learn how to use, so the use of AI in this course is allowed with some restrictions.

What is allowed: AI is allowed, and expected, to be used in the technical aspects of this course, including system construction for student projects.

Mandatory citation: all use of generative AI, whether in technical or writing assignments, must be clearly documented and properly cited.

What is prohibited with some exceptions: the use of generative AI for writing assignments. All writing must be your own. Having AI complete a writing assignment on your behalf will be treated similarly to having a peer write the assignment for you. However, there are some acceptable uses of involving AI in writing assignments, similar to acceptable uses of involving a peer in writing assignments: brainstorming ideas, identifying gaps or errors in your drafts, and proofreading.

Essentially, your writing should be your own, but just like how it is sometimes okay to ask others for brainstorming or proofreading help, you can use generative AI to brainstorm or clean things up.

All other uses of generative AI in this course are prohibited.

Consequences of misuse: misuse of AI, including improperly citing or documenting use, may result in a reduction in assignment grades, and in severe cases will be treated as plagiarism or academic dishonesty.

Academic Standards

Academic Standards exist to promote authentic scholarship, support the institution's goal of maintaining high standards of academic excellence, and encourage continued ethical behavior of faculty and students to cultivate an educational community which values integrity and produces graduates who carry this commitment forward into professional practice.

As members of the George Mason University community, we are committed to fostering an environment of trust, respect, and scholarly excellence. Our academic standards are the foundation of

this commitment, guiding our behavior and interactions within this academic community. The practices for implementing these standards adapt to modern practices, disciplinary contexts, and technological advancements. Our standards are embodied in our courses, policies, and scholarship, and are upheld in the following principles:

- **Honesty:** Providing accurate information in all academic endeavors, including communications, assignments, and examinations.
- **Acknowledgement:** Giving proper credit for all contributions to one's work. This involves the use of accurate citations and references for any ideas, words, or materials created by others in the style appropriate to the discipline. It also includes acknowledging shared authorship in group projects, co-authored pieces, and project reports.
- **Uniqueness of Work:** Ensuring that all submitted work is the result of one's own effort and is original, including free from self-plagiarism. This principle extends to written assignments, code, presentations, exams, and all other forms of academic work.

Violations of these standards—including but not limited to plagiarism, fabrication, and cheating—are taken seriously and will be addressed in accordance with university policies. The process for reporting, investigating, and adjudicating violations is [outlined in the university's procedures](#). Consequences of violations may include academic sanctions, disciplinary actions, and other measures necessary to uphold the integrity of our academic community.

The principles outlined in these academic standards reflect our collective commitment to upholding the highest standards of honesty, acknowledgement, and uniqueness of work. By adhering to these principles, we ensure the continued excellence and integrity of George Mason University's academic community.

Student responsibility: Students are responsible for understanding how these general expectations regarding academic standards apply to each course, assignment, or exam they participate in; students should ask their instructor for clarification on any aspect that is not clear to them.

Accommodations for Students with Disabilities

Disability Services at George Mason University is committed to upholding the letter and spirit of the laws that ensure equal treatment of people with disabilities. Under the administration of University Life, Disability Services implements and coordinates reasonable accommodations and disability-related services that afford equal access to university programs and activities. Students can begin the registration process with Disability Services at any time during their enrollment at George Mason University. If you are seeking accommodations, please visit <https://ds.gmu.edu/> for detailed information about the Disability Services registration process. Disability Services is located in Student Union Building I (SUB I), Suite 2500. Email: ods@gmu.edu. Phone: (703) 993-2474.

Student responsibility: Students are responsible for registering with Disability Services and communicating about their approved accommodations with their instructor *in advance* of any relevant class meeting, assignment, or exam.

FERPA and Use of GMU Email Addresses for Course Communication

The [Family Educational Rights and Privacy Act \(FERPA\)](#) governs the disclosure of [education records for eligible students](#) and is an essential aspect of any course. **Students must use their GMU email account** to receive important University information, including communications related to this class. Instructors

will not respond to messages sent from or send messages regarding course content to a non-GMU email address.

Student responsibility: Students are responsible for checking their GMU email regularly for course-related information, and/or ensuring that GMU email messages are forwarded to an account they do check.

Title IX Resources and Required Reporting

As a part of George Mason University's commitment to providing a safe and non-discriminatory learning, living, and working environment for all members of the University community, the University does not discriminate on the basis of sex or gender in any of its education or employment programs and activities. Accordingly, **all non-confidential employees, including your faculty member, have a legal requirement to report to the Title IX Coordinator, all relevant details obtained directly or indirectly about any incident of Prohibited Conduct** (such as sexual harassment, sexual assault, gender-based stalking, dating/domestic violence). Upon notifying the Title IX Coordinator of possible Prohibited Conduct, the Title IX Coordinator will assess the report and determine if outreach is required. If outreach is required, the individual the report is about (the "Complainant") will receive a communication, likely in the form of an email, offering that person the option to meet with a representative of the Title IX office. For more information about non-confidential employees, resources, and Prohibited Conduct, please see [University Policy 1202](#): Sexual and Gender-Based Misconduct and Other Forms of Interpersonal Violence. Questions regarding Title IX can be directed to the Title IX Coordinator via email to TitleIX@gmu.edu, by phone at 703-993-8730, or in person on the Fairfax campus in Aquia 373.

Student opportunity: If you prefer to speak to someone *confidentially*, please contact one of Mason's confidential employees in Student Support and Advocacy ([SSAC](#)), Counseling and Psychological Services ([CAPS](#)), Student Health Services ([SHS](#)), and/or the [Office of the University Ombudsperson](#).