



Course Syllabus

T – 81-559 Applications of Generative AI and Large Language Models, Fall 2024

The Henry Edwin Sever Institute | James McKelvey School of Engineering

Washington University in St. Louis

COURSE OVERVIEW

This course covers the dynamic world of Generative Artificial Intelligence providing hands-on practical applications of Large Language Models (LLMs) and advanced text-to-image networks. Using Python as the primary tool, students will interact with OpenAI's models for both text and images. The course begins with a solid foundation in generative AI principles, moving swiftly into the utilization of LangChain for model-agnostic access and the management of prompts, indexes, chains, and agents. A significant focus is placed on the integration of the Retrieval-Augmented Generation (RAG) model with graph databases, unlocking new possibilities in AI applications. As the course progresses, students will delve into sophisticated image generation and augmentation techniques, including LoRA (Low-Rank Adaptation), and learn the art of fine-tuning generative neural networks for specific needs. The final part of the course is dedicated to mastering prompt engineering, a critical skill for optimizing the efficiency and creativity of AI outputs. Ideal for students, researchers, and professionals in computer science or related fields, this course offers a transformative learning experience where technology meets creativity, paving the way for innovative applications in the realm of Generative AI. Note: This course will require the purchase of up to \$100 in OpenAI API credits to complete the course.

- This is a technical course that will require programming in the Python programming language. It is not necessary to know Python prior to this course; however, familiarity of at least one programming language is assumed.

COURSE SECTIONS

- **Section 1 (hybrid): room Lopata Hall / 202, 6:00, 8/27/2024, 9/24/2024, 10/22/2024, 11/26/2024**

For hybrid students, this course will meet for four in-person sessions at the WUSTL Danforth Campus. Your attendance at these three in-person meetings is encouraged. I will Zoom-broadcast the four class sessions and post a recording the following day. Additionally, I conduct a weekly Zoom-only sync meeting for each week that we do not meet in person. These weekly meetings supplement the prerecorded material in Canvas for each module. I record and post all sessions in Canvas. Generally, the weekly sync meetings will last one hour or less. You can find all Zoom meetings in Canvas. There are also weekly prerecorded lessons for each week.

INSTRUCTOR OVERVIEW



Instructor: Jeff Heaton
Vice President, Reinsurance Group of America (RGA)
Adjunct Instructor, Washington University (WUSTL)

Email: jheaton@wustl.edu

Phone: (636) 525-1842

Note: email or MS-Teams is the preferred means of contact for me.

Office Hours: By appointment

Response Time: I will respond to emails within 24-48 hours.

I started my career in Information Technology (IT), working as a computer programmer in languages such as C/C++, Java, Python, SQL, and PHP. Backend systems that perform complex calculations on high-performance computing (HPC) operations for financial forecasting have always been my specialty. In my early career, I've worked for Anheuser Busch, Monsanto, Boeing, and MasterCard. For the past 21 years, I've worked for Reinsurance Group of America (RGA), and am currently Vice President of Data Science at RGA.

I've always been passionate about sharing my knowledge and interest in machine learning with others. I run a popular YouTube Channel with over 70K subscribers and have over 10K Twitter followers. Teaching deep learning for Washington University was a natural extension of this, as I enjoy working with students on the latest deep learning technologies. I am often able to recruit interns from my classes to work with these technologies in my group at RGA.

I hold a Masters of Information (MIM) from Washington University, a Ph.D. in Computer Science, and am a senior member of IEEE. I am the author of a couple of books on AI and a few peer-reviewed papers for the Journal of Machine Learning Research and Genetic Programming and Evolvable Machines.

LEARNING OBJECTIVES

By the end of this course, you will be able to:

- Explain the differences between the major technologies of generative AI.
- Design a GenAI based solution to a Kaggle competition
- Measure the performance of a GenAI-based solution.
- Apply LangChain to large language models.
- Check and evaluate your Python code and resolve issues.



CLASS MEETING

Sync Sessions

These sessions will be held through Zoom, all online this semester. Their times will be announced the first week of class, I will define their time some to accommodate students outside the university's time zone. I try to do these once a week usually on a Tuesday. I will always record the sync sessions and place them in Canvas. You'll get a Canvas announcement to let you know when it will be held. They will consist of about 15-20 minutes of talking about the week's content. Then, there will be plenty of time to answer or discuss any questions that you may have about assignments, content or anything else. This is not a time when new content will be introduced, but instead, you will have time to gain clarity about content for the week or your assignments

Participation and Attendance

This course will not meet at a particular time each week. All course goals, session learning objectives, and assessments are supported through classroom elements that can be accessed at any time. To measure class participation (or attendance), your participation in threaded discussion boards is required, graded, and paramount to your success in this course. Please note that any scheduled synchronous meetings are optional. While your attendance is highly encouraged, it is not required and you will not be graded on your attendance or participation. It is important, though, that you post to the discussions in a timely manner. Do not post ahead of schedule. While it is acceptable (and encouraged) for students to work ahead of the stated schedule, work should be posted only during the week it is due. In other words, students should not post discussion contributions in Session 6 during Session 4. If students wish to compose their assignments in advance, they should maintain those on their personal computer and post them only during the week they are due.

CLASS TEXTS / MATERIALS / TOOLS

The reading material for this course is contained in the following freely available PDF and also at my GitHub repository. Readings in both the PDF and GitHub are organized into the same 14 modules as this course.

- https://github.com/jeffheaton/app_generative_ai



GRADE COMPOSITION

Major Coursework Components	Component Proportion	Coursework Sub-component	Sub-component Proportion	Points
Icebreaker Introduction	5%			5
Programming Assignments	50%	Program 1	5 pts	50
		Program 2	5 pts	
		Program 3	5 pts	
		Program 4	5 pts	
		Program 5	5 pts	
		Program 6	5 pts	
		Program 7	5 pts	
		Program 8	5 pts	
		Program 9	5 pts	
		Program 10	5 pts	
Kaggle Competition (team)	25%	Kaggle Submission	10 pts	25
		Post Solution	15 pts	
Final Project (team)	20%	Project	15	20
		Individual Contribution	5	

Total Point: 100

COURSEWORK COMPONENTS DESCRIPTIONS

Refer to the Canvas assignment rubrics for a complete description of each assignment, due dates in Canvas are the official due dates for all assignments.

Ice Breaker Introductions – This assignment allows the students to get to know each other. I will post an ice breaker for myself to start the process. Please include, your education background, your current field of study, what you hope to get from this course, and a fun activity or interesting fact about yourself.

Group Selection – The Kaggle competition and final projects for this semester can be completed as a group project. For the maximum group size, refer to the Kaggle assignment in Canvas. Students will be allowed to form their own group, be assigned a group by me, or complete the assignment individually. Once groups are formed, they must submit a plan that includes a group leader selection. Group members will provide me with feedback on individual member participation, which will be reflected in your grade.



Programming Assignments – There are ten weekly programming assignments that will check your knowledge of each module's content. Five points are possible. 1 point for completed submission, 2 for program correctness, 2 points for correct output.

Kaggle Competition – Each semester I create a new Kaggle competition for the class. The competition is hosted on Kaggle In-Class. Teams will create a solution and submit their solutions. Kaggle will rank the solutions from most accurate to least. Your grade is not dependent on your ranking.

Final Project – Generative AI is a rapidly evolving field. For this assignment you will read an academic paper dealing with an advanced GenAI technique. You will answer a series of questions on this paper in a written report. The same teams from the Kaggle competition will complete this project.

UNIVERSITY COURSE EVALUATIONS

Note to the student: Each student is asked to thoughtfully complete and submit the university course evaluation. It is fully anonymous and is one of the best tools we have to continue to improve the student experience for all students. Course evaluations become available 5 days before the end of the course and remain open for 8 days. They can be accessed through WebStac.

GRADING POLICIES

Late Work Policy

Late submissions of any assignment will reduce the assignment grade by 2 points per day following the due date/time. Exceptions may include a mutually agreed arrangement made in advance or bona fide emergencies (be prepared to supply documentation). Problems with your personal computer (such as computer crashes) or sudden loss of Internet access at your home do not count as an emergency. Please let me know if something significant happens to you during the term that could interfere with your submitting class assignments on time (death in family, loss of job, etc.) Submitting your work late may be a calculated decision on your part. I certainly understand this. However, if you plan to or choose to submit an assignment late, please let me know in advance of the due date. Late submission of course assignments should be the exception and will be penalized as described herein. Under no circumstances should students submit an assignment for credit after the last day of final exams as identified in the official University academic calendar for the applicable quarter. **Any assignment submitted more than a week late will not receive any credit.**



Grading Scale:

Letter Grade	%	Points Toward GPA	Letter Grade	%	Points Toward GPA
A+	$\geq 97\%$	4.0	C+	77% - 79%	2.3
A	93% - 97%	4.0	C	73% - 76%	2.0
A-	90% - 92%	3.7	C-	70% - 72%	1.7
B+	87% - 89%	3.3	D+	67% - 69%	1.3
B	83% - 86%	3.0	D	65% - 66%	1.0
B-	80% - 82%	2.7	F	< 65%	0.0

COURSE SCHEDULE

Module	Content
<u>Module 1</u> Meet on 08/27/2024	Module 1: Introduction to Generative AI <ul style="list-style-type: none">1.1: Course Overview1.2: Generative AI Overview1.3: Introduction to OpenAI1.4: Introduction to LangChain1.5: Prompt EngineeringWe will meet on campus this week! (first meeting)
<u>Module 2</u> Week of 09/03/2024	Module 2: Prompt Based Development <ul style="list-style-type: none">2.1: Prompting for Code Generation2.2: Handling Revision Prompts2.3: Using a LLM to Help Debug2.4: Tracking Prompts in Software Development2.5: Limits of LLM Code Generation<u>Module 1 Program</u> due: 09/04/2024Icebreaker due: 09/04/2024
<u>Module 3</u> Week of 09/10/2024	Module 3: Introduction to Large Language Models <ul style="list-style-type: none">3.1: Foundation Models3.2: Text Generation



Module	Content
	<ul style="list-style-type: none">• 3.3: Text Summarization• 3.4: Text Classification• 3.5 LLM Writes a Book• Module 2 Program due: 09/11/2024
Module 4 Week of 09/17/2024	Module 4: LangChain: Chat and Memory <ul style="list-style-type: none">• 4.1: LangChain Conversations• 4.2: Conversation Buffer Window Memory• 4.3: Conversation Token Buffer Memory• 4.4: Conversation Summary Memory• 4.5: Persisting Langchain Memory• Module 3: Program due: 09/18/2024
Module 5 Meet on 09/24/2024	Module 5: LangChain: Data Extraction <ul style="list-style-type: none">• 5.1: Structured Output Parser• 5.2: Other Parsers (CSV, JSON, Pandas, Datetime)• 5.3: Pydantic parser• 5.4: Custom Output Parser• 5.5: Output-Fixing Parser• Module 4 Program due: 09/25/2024• We will meet on campus this week! (second meeting)
Module 6 Week of 10/01/2024	Module 6: Retrieval-Augmented Generation (RAG) <ul style="list-style-type: none">• 6.1 Introduction to RAG• 6.2 Introduction to ChromaDB• 6.3 Understanding Embeddings• 6.4 Q&A Over Documents• 6.5 Embedding Databases• Module 5 Program due: 10/02/2024



Module	Content
<u>Module 7</u> Week of 10/15/2024	Module 7: LangChain: Agents <ul style="list-style-type: none">7.1: Introduction to LangChain Agents7.2: Understanding LangChain Agent Tools7.3: LangChain Retrieval and Search Tools7.4: Constructing LangChain Agents7.5: Custom Agents <p>Module 6 Program due: 10/16/2024</p>
<u>Module 8</u> Meet on 10/22/2024	Module 8: Kaggle Assignment <ul style="list-style-type: none">8.1: Introduction to Kaggle8.2: Kaggle Notebooks8.3: Small Large Language Models8.4: Accessing Small LLM from Kaggle8.5: Current Semester's Kaggle <p>Module 7 Program due: 10/23/2024</p> <p>We will meet on campus this week! (third meeting)</p>
<u>Module 9</u> Week of 10/29/2024	Module 9: Introduction to StreamLit <ul style="list-style-type: none">9.1: Running StreamLit in Google Colab9.2: StreamLit Introduction9.3: Understanding Streamlit State9.4: Creating a Chat Application9.5: More Advanced Chat Application <p>Module 8 Program due: 10/23/2024</p>
<u>Module 10</u> Week of 11/05/2024	Module 10: Text to Image Generative AI <ul style="list-style-type: none">10.1: Stable Diffusion10.2: Calling Dall-E10.3: Upscaling Images10.4: Introduction to Dreambooth



Module	Content
	<ul style="list-style-type: none">• 10.5: Adding Images to ChatBot• Module 9 Program due: 11/06/2024
Module 11 Week of 11/12/2024	Module 11: Fine Tuning <ul style="list-style-type: none">• 11.1: When is fine tuning necessary• 11.2: Preparing a dataset for fine tuning• 11.3: OpenAI Fine Tuning• 11.4: Application of Fine Tuning• 11.5: Evaluating Fine Tuning and Optimization• Module 10 Program due: 11/13/2024
Module 12 Week of 11/19/2024	Module 12: Prompt Engineering <ul style="list-style-type: none">• Kaggle Assignment due: 04/15/2024 (approx 4-6PM, due to Kaggle GMT timezone)• 12.1 Towards the Perfect Prompt• 12.2 Prompt Engineering for Images• 12.3 Prompt Engineering for LLM• 12.4 Advanced Tips for Prompt Engineering• 12.5 Evaluating Bias in Prompts
Module 13 Meet on 11/26/2024	Module 13: Speech Processing <ul style="list-style-type: none">• 13.1: Audio Generation and Recognition• 13.2: OpenAI Speech Generation• 13.3: OpenAI Speech Recognition• Part 13.4: A Voice-Based ChatBot• 13.5: Future Directions in GenAI• We will meet on campus this week! (fourth meeting)• Final project due: 12/03/2024



I. **POLICIES**

Use of Laptop Computers and Electronic Devices in the Classroom

Laptop & tablet computers, smart phones and other electronic devices can be helpful in taking notes, providing tools for course exercises and referencing course related materials. However, they can also be distracting when used for non-course related activities such as emailing & texting, posting on social media, reading news sites, shopping online, or looking at YouTube videos. Some students have even been observed working on class assignments for the same or other courses. As common sense suggests, and a March 2013 study by Faria Sana, Tina Weston and Nicholas J. Cepeda confirmed, students who are multitasking during class have less understanding and recall of what's being discussed. The study also found that "participants who were in direct view of a multitasking peer scored lower on a test compared with those who were not."

As mentioned earlier this course is part of a professional, graduate program. Consequently, it is expected that students conduct themselves in a professional manner. This includes being engaged in the class proceedings, by attentive listening, critical thinking, asking appropriate questions and participating in active discussion. Your attendance and participation in class is important for the class and is expected to be more than just physical attendance. Engaging in non-class related activities during class time is not acceptable and disrespectful of the lecturer and other students.

Privacy and Security

Recording of class sessions either audio or video is prohibited without permission from the instructor and the other class members. The instructor will make recordings of this course available through Canvas.

Collaboration:

With the exception of your team projects, all assignments are to be completed on your own. You are encouraged to discuss ideas and techniques broadly with other class members, but all written or presentation work, whether in preliminary or final form, is to be generated by you working alone. If in doubt - *ask*.

Language Sensitivity

When in the classroom, all students should speak English at all times. While meeting with classmates on a classroom project, speak a language that every student present (in your group) understands, without exception.

Professionalism:

You are part of a professional, graduate program. Consequently, it is expected that your fellow students conduct yourselves in a *professional* manner. This includes being on time for classes and meetings, being prepared, and participating in class discussions, group activities, projects, etc. The level of professionalism you exhibit throughout the course will impact your final grade. It directly affects the participation portion of the grade but is also taken into consideration in all other aspects of the course as it reflects the overall quality of professional performance.



II. SEVER/UNIVERSITY POLICIES

Ethics of Academic Integrity (SEAS)

All students in the School of Engineering & Applied Science are expected to conform to high standards of conduct. This statement on student academic integrity is intended to provide guidelines on academic behaviors which are not acceptable.

Engineering courses typically have many problem sets assigned as homework. You are not allowed to collaborate when solving homework problems, performing lab experiments, writing or documenting computer programs, or writing reports unless the instructor specifically states otherwise.

It is dishonest and a violation of academic integrity if:

1. You turn in work which is represented as yours when in fact you have significant outside help. When you turn in work with your name on it, you are in effect stating that the work is yours, and only yours.
2. You use the results of another person's work (exam, homework, computer code, lab report) and represent it as your own, regardless of the circumstances.
3. You request special consideration from an instructor when the request is based upon false information or deception.
4. You submit the same academic work to two or more courses without the permission of each of the course instructors. This includes submitting the same work if the same course is retaken.
5. You willfully damage the efforts of other students.
6. You use prepared materials in writing an in-class exam except as approved by the instructor.
7. You write on or make erasures on any test material or class assignment being submitted for re-grading.
8. You collaborate with other students planning or engaged in any form of academic dishonesty.
9. You turn in work, which is represented as a cooperative effort, when in fact you did not contribute your fair share of the effort.
10. You do not use proper methods of documentation. For example, you should enclose borrowed information in quotation marks; acknowledge material that you have abstracted, paraphrased or summarized; cite the source of such material by listing the author, title of work, publication, and page reference.



III. WASHINGTON UNIVERSITY IN ST LOUIS SUPPLEMENTAL RESOURCES

1. **Disability Resources**: If you have a disability that requires an accommodation, please speak with instructor and consult the **Disability Resource Center** at Cornerstone (cornerstone.wustl.edu/). Cornerstone staff will determine appropriate accommodations and will work with your instructor to make sure these are available to you.
2. **English writing support**: For additional help on your writing, consult the expert staff of **The Writing Center** (writingcenter.wustl.edu) in Olin Library (first floor). It can be enormously helpful to ask someone outside a course to read your essays and to provide feedback on strength of argument, clarity, organization, etc.

The **Engineering Communication Center** (<http://engineering.wustl.edu/current-students/student-services/Pages/default.aspx>) offers students in the School of Engineering and Applied Sciences help with oral presentations, writing assignments, and other communications projects, as well as job-search documents such as resumes and cover letters

3. **English competence**: Students are encouraged to check their grammar and spelling before submitting their written works. Although, students are free to choose whatever tools best fit their need, some of the common tools for grammar, spelling, and citing references can be found in the list below.
 - a. <https://www.merriam-webster.com> – Merriam-Webster Dictionary [Free]
 - b. <https://www.grammarly.com> – Grammarly [Free & Paid Service]
 - c. <http://www.gingersoftware.com> – Ginger [Free & Paid Service]
 - d. <http://www.citationmachine.net> – Citation Machine [Free & Paid Service]
4. **Bias reporting**: The University has a process through which students, faculty, staff and community members who have experienced or witnessed incidents of bias, prejudice or discrimination against a student can report their experiences to the University's Bias Report and Support System (BRSS) team. See: brss.wustl.edu/
5. **Mental health service**: Mental Health Services' professional staff members work with students to resolve personal and interpersonal difficulties, many of which can affect the academic experience. These include conflicts with or worry about friends or family, concerns about eating or drinking patterns, and feelings of anxiety and depression. See: shs.wustl.edu/MentalHealth
6. **Sexual Harassment**: Sexual harassment is a form of discrimination that violates university policy and will not be tolerated. It is also illegal under state and federal law. Title IX of the Education Amendments of 1972 prohibits discrimination based on sex (including sexual

harassment and sexual violence) in the university's educational programs and activities. Title IX also prohibits retaliation for asserting claims of sex discrimination. The university has designated the Title IX Coordinator identified below to coordinate its compliance with and response to inquiries concerning Title IX.

For more information or to report a violation under the Policy on Discrimination and Harassment, please contact:

Discrimination and Harassment Response Coordinators

Apryle Cotton, Asst. Vice Chancellor for Human Resources
Section 504 Coordinator
Phone: 314-362-6774
Email: apryle.cotton@wustl.edu
Leanne Stewart, Employee Relations Manager
Phone: 314-362-8278
Email: leannerstewart@wustl.edu

Title IX Coordinator

Jessica Kennedy, Director of Title IX Office
Title IX Coordinator
Phone: 314-935-3118
Email: jwkennedy@wustl.edu

You may also submit inquiries or a complaint regarding civil rights to the United States Department of Education's Office of Civil Rights at 400 Maryland Avenue, SW, Washington, DC 20202-1100 or by visiting the [U.S. Department of Education website](#) or calling 800-421-3481.

IV. Technical Requirements

1. Students must have access to a reliable Internet connection that can support Zoom.
2. Students must be able to execute and run Python programs. Instructions will be provided for setting up a Python/TensorFlow environment either on the student's own computer or through Google CoLab
3. For technical support, connecting to Zoom or the University Computer Network, please contact the WUSTL HelpDesk, <https://one.wustl.edu/task/all/itsupport>