Course Information

CS106EA: Exploring Artificial Intelligence

Patrick Young, Stanford University
Winter 2026

Welcome to CS106EA. This course will provide a broad overview of Artificial Intelligence, with an emphasis on Neural-Network-based systems. We'll be focusing on how AI works at a conceptual level while also providing a good amount of hands-on practice with AI systems using Google Colab, a platform widely used by AI researchers and data scientists.

Course Goals

Our objective is for students completing CS106EA to:

- Develop a strong understanding of how AI systems work.
- Gain familiarity with key concepts and terminology used in the field of AI.
- Be well prepared for interdisciplinary AI classes at Stanford taught outside of Computer Science, with applications such as AI and healthcare, social impacts of AI, business applications of AI, and more.
- Be able to confidently engage with AI-related articles, whether written for a general audience (e.g., The New York Times, Bloomberg, Scientific American) or a technical but non-specialized audience (e.g., IEEE Spectrum, Communications of the ACM).
- Be prepared as AI continues to expand in fields such as biology, chemistry, economics, and the social sciences—CS106EA will provide a solid foundation to engage with AI research in your own field.
- Have enhanced readiness for work as a Product Manager on an AI team by gaining a deeper understanding of AI systems and their capabilities.
- Be able to critically evaluate the strengths, limitations, and ethical implications of AI applications.
- Have confidence engaging with AI topics, whether through hands-on work, strategic decision-making, or policy discussions.

For a detailed breakdown of the specific topics covered in this course, please refer to the list later in this handout.

Course Staff

All e-mail addresses listed in this section are @stanford.edu.

Lecturer

Dr. Patrick Young
Office Hours: Check Canvas
Email: patrick.young @

Teaching Assistants

Aanika Kaur Atluri, aatluri @
Yuliia Murakami, ym256 @

The fastest way to contact the Teaching Staff is to post on our Ed Discussion board which can be found on Canvas. All of us will be monitoring. If you want to ask something privately, just put up a private Ed Discussion post.

Office hours are subject to change. For the latest information on office hours, including both times and locations, please check Canvas.

Prerequisites

The official prerequisite for the course is CS106A. Students who have not completed CS106A but have a basic understanding of programming are welcome to take the course. You will be running AI programs in our class but not writing them.

Students taking the course should have:

- a basic understanding of how a program works,
- acquaintance with basic data types used in programming (integers, floating point, and strings),
- familiarity with basic control structures (if-statements, for-loops).

While AI programs are generally written in Python, the above level of familiarity with any programming language should be sufficient.

There is no Math prerequisite. Modern AI does depend on Statistics, and we do assume that all Stanford students have at least a basic understanding of Mean (i.e., Average) and a general sense of what Standard Deviation and Variance represent.

Grading

Grades will be divided evenly between exams and homeworks.

Midterm	25%
Final	25%
Assignments	50%

Grades on Canvas

We will be posting your grades on the Canvas website (see below). Please make sure you check your grades as soon as they are posted. Any missing grades or grading errors must be reported within one week of the posting date. *All grades will be considered final, one week after they are posted.*

Homeworks

There will be a weekly homework assignment which will be released each Thursday and due the following Wednesday night at 11:59pm. We'll be using Gradescope for submission (see Gradescope tab/section on Canvas). Assignments are designed to reinforce concepts taught in lecture. A variety of different methods will be used including:

- Running AI systems on Google Colab; analyzing and reporting the results; and considering what the implications are, tying them back into concepts taught in lecture.
- Examining AI systems that you may have experience with in real life (e.g. ChatGPT, image classification on Android or Apple devices, music or video recommendation systems, translation software) and reflecting on how they might be developed based on concepts taught in lecture.

Using AI

As a class on AI, we encourage you to experiment with AI including when doing your course work for this class. However, please remember that these rules are specific to CS106EA. Other classes at Stanford may have more restrictive policies on AI usage that reflect their own teaching objectives, so always check the guidelines for each course.

For the homework, you may discuss and refine your answers using AI systems. However, the final submission should reflect your own thoughts written in your own words. **Do not copy and paste AI-generated content directly into your homework submissions.** This approach will help you engage more deeply with the material and remember it better. On the exams you will need to remember, understand, and synthesize the ideas from the homework on your own.

In addition, keep in mind that AI systems are not 100% reliable and are known to sometimes produce complete fabrications, often referred to as hallucinations. As we'll learn later in the quarter, Large Language Models such as ChatGPT are trained to give plausible responses, not necessarily truthful ones, so they are very good at coming up with responses that seem plausible but are false. If you run into a situation where you aren't convinced the AI is correct, there's a good chance you are right, and the AI is wrong. When in doubt you can always post to our Ed Discussion board.

The exams are closed book, closed note, closed computer, so you won't have access to AI during exams. However, you should definitely feel free to feed course material such as the lecture notes, review handouts, and homework handouts into AI systems and ask them to quiz you on how well you know the material.

Working with Others

We'll use rules similar to our AI usage rules. You are welcome to discuss your homework with other students. In fact, a good group discussion is an excellent way to engage with the material and may lead to better memory recall than studying by yourself. However, as with AI, we do ask that you ultimately **write your own answers from scratch**. Do not copy and paste from other students.

For our Google Colab exercises, you are welcome to observe others running them and discuss results, **but you must also rerun them on your own**. Working through the process yourself, considering what happens at each step, and evaluating your results will improve your understanding and recall.

Keep in mind that most AI processes are non-deterministic, meaning there is an element of randomness, and they typically will not produce the same results even under identical conditions. When we ask you to take a screenshot or list specific results from a program, **your program's outputs should be subtly (and in some cases, not so subtly) different from your classmates**. In fact, comparing results with your fellow classmates and discussing the reasons why results differ may give you additional insight into how these systems work.

Late Policy

Assignments turned in late will be penalized 10% for each day that has passed since the original due date. *No assignment may be turned in more than a week after its original due date and no assignments will be accepted after 11:59pm on the evening of Friday March 13th*, as we'll need the time to get grading completed before Stanford's final grading deadline.

I realize that you do have other classes and other responsibilities. Therefore, you will be given a late allowance of three late days that can be used to excuse late assignments. Simply submit your assignment late on Canvas, no need to notify us. *You are expected to use these for any problems that come up, including getting sick, or other emergencies*.

This allowance may be used for a single assignment, or it may be divided for use on multiple assignments. For example, if you turn in one assignment three days late, you've just used up your entire late allowance. However, if instead you turn in the assignment two days late you still have an additional late day which you can use for another assignment.

Exams

Exams reinforce the course material by encouraging students to review the concepts and terms presented in the class. The process of studying for exams strengthens understanding, reinforces key ideas, and helps students retain information more effectively.

We will release review questions and lists of key terms each lecture alongside the lecture slides. These resources are intended to highlight the most important ideas covered and help focus your attention on core concepts and terms.

Exam questions will draw from the review material provided as well as from the concepts covered in the homework assignments. If you have a good understanding of the review questions, master the listed terms, and comprehend the main points of the homework assignments you should be well prepared for the exams.

Exams will be closed book, closed note, closed computer. We will have two exams:

- **Midterm Exam**: Tuesday, February 10th during our regular lecture time (location TBD).
- **Final Exam**: Monday, March 16th from 9:30-11:30am.

Important notes about the Final Exam:

1) Our final will only be two hours long, so we will start at 9:30am.
2) The CS105 and CS106 classes have a special exam time slot which is not based on their regular class meeting time, we will be using this special exam time slot.

Course Material

For those who want a better idea of what we'll actually be covering, here is a high-level overview of the main course topics along with estimated lecture counts to give you an idea

of how much time will be spent on each section. *If these topics all seem a bit mysterious to you, you are in the right class*.

Introduction to AI (2 lectures)

Provides an overview of AI. Introduces Machine Learning and outlines the Machine Learning process. Introduces key AI terms and concepts that will appear throughout the quarter.

Neural Networks (3 lectures)

This lecture series provides a thorough grounding on Neural Networks. This is the technology underpinning the vast majority of advanced AI systems today. We'll study how these systems work, how they are trained, and what happens when the training goes wrong.

Real-World Considerations (2 lectures)

These lectures focus on practical issues in developing and deploying AI systems. We also study the types of problems modern AI systems excel at and those that are inherently challenging or may be fundamentally unsolvable with neural-network-based systems.

Specialized AI Systems (3 lectures)

While basic neural networks can work for some problems, specialized architectures are often required to handle specific types of data effectively. In this section, we explore neural networks designed for image processing and natural language understanding (e.g., English). These architectures not only show how we can address unique challenges found in images and natural language but also serve as exemplars of how AI researchers adapt models to work with diverse and complex inputs.

Transformers and Large Language Models (4 lectures)

Transformers are a groundbreaking architecture that powers Large Language Models (LLMs) such as ChatGPT. We begin by examining the Transformer Architecture on its own, focusing on features that make it powerful and versatile. As we will see these features have enabled their use far beyond LLMs. We then take a look at LLMs themselves, studying how they work and how they are trained.

Advanced Topics (2 lectures)

At this stage, you will have a solid grounding in major AI concepts. These lectures will explore topics to be determined after consulting with the class and the teaching team. Potential topics may include autonomous vehicles and image generation systems.

Social Impacts and Future Directions (2 lectures)

In this series we consider societal, ethical, and economic issues that AI use may pose in the short-, medium-, and long-term. We also explore active research areas and discuss potential future directions of AI research.

Schedule

Reminder: homeworks are not shown on the schedule but **assignments will go out each Thursday lecture and will be due each Wednesday at 11:59pm**.

Date	Week	Tuesday	Thursday
1/6	1	Introduction to AI, Introduction to Machine Learning	Machine Learning (cont.), Classical Machine Learning, Google Colab
1/13	2	Introduction to Neural Networks	Neural Networks: Key Components & Training Overview
1/20	3	Neural Networks: Training Details, Issues and Solutions	Real World Considerations: Development and Deployment
1/27	4	Real World Considerations: What can AI Predict? Poor Applications of AI	Specialized AI Systems for Image Processing
2/3	5	Specialized Techniques for Natural Language Processing (NLP)	Specialized AI Systems for Natural Language Processing and Time Series Data
2/10	6	**Midterm** Location TBD	Why Transformers? Key Features of Transformers
2/17	7	The Transformer Architecture	Full Models with Transformers
2/24	8	Details on Large Language Models	Advanced Topics
3/3	9	Advanced Topics	Social Impacts
3/10	10	Social Impacts and Future Directions	TBD

Course Website

We will be using Canvas for our class website. You can access Canvas at:

http://canvas.stanford.edu/

Incompletes

If you have a serious medical or family emergency and cannot complete the work in this course, you may contact me to request an incomplete. I reserve incompletes only for emergencies, so I do not grant incomplete grades for poor performance on the assignments or exams, nor do I offer incompletes for busy work schedules.

In order to be eligible for an incomplete, you must have completed all of the assignments (with the possible exception of the most-recently-due assignment) and must have shown satisfactory academic progress in the class.

Office of Accessible Education

Students who may need an academic accommodation based on the impact of a disability must initiate the request with the Student Disability Resource Center (SDRC) located within the Office of Accessible Education (OAE). SDRC staff will evaluate the request with required documentation, recommend reasonable accommodations, and prepare an Accommodation Letter for faculty dated in the current quarter in which the request is being made. Students should contact the SDRC as soon as possible since timely notice is needed to coordinate accommodations. The OAE is located at 563 Salvatierra Walk (phone: 650-723-1066)

Financial Aid

The School of Engineering has asked that we pass the following information on to all students on Financial Aid:

All students should retain receipts for books and other course-related expenses, as these may be qualified educational expenses for tax purposes. If you are an undergraduate receiving financial aid, you may be eligible for additional financial aid for required books and course materials if these expenses exceed the aid amount in your award letter. For more information, review your award letter or visit the Student Budget website:

https://financialaid.stanford.edu/undergrad/budget/index.html

Contingency Planning

Stanford as an institution is committed to the highest quality education, and as your teaching team, our first priority is to uphold your educational experience. To that end we are committed to following the syllabus as written here, including through short- or long-

term disruptions, such as public health emergencies, natural disasters, or protests and demonstrations. However, there may be extenuating circumstances that necessitate some changes. Should adjustments be necessary, we will communicate clearly and promptly to ensure you understand the expectations and are positioned for successful learning.

The Stanford Honor Code

The Honor Code is an undertaking of the Stanford academic community, individually and collectively. Its purpose is to uphold a culture of academic honesty.

Students will support this culture of academic honesty by neither giving nor accepting unpermitted academic aid in any work that serves as a component of grading or evaluation, including assignments, examinations, and research.

Instructors will support this culture of academic honesty by providing clear guidance, both in their course syllabi and in response to student questions, on what constitutes permitted and unpermitted aid. Instructors will also not take unusual or unreasonable precautions to prevent academic dishonesty.

Students and instructors will also cultivate an environment conducive to academic integrity. While instructors alone set academic requirements, the Honor Code is a community undertaking that requires students and instructors to work together to ensure conditions that support academic integrity.