

AI Applications & Ethics

Course Outline & Objectives

Course Description:

In AI Applications & Ethics, students develop AI literacy while exploring how modern artificial intelligence systems are designed, used, and evaluated. Students investigate how AI models generate and retrieve information, analyze real-world applications, and examine issues such as accuracy, bias, privacy, intellectual property, and responsible use. They apply prompt design strategies, evaluate AI-generated content, conduct investigations into model behavior, and communicate findings through research and project-based work. Throughout the course, students explore the ethical, societal, and technological impacts of AI while developing critical thinking, problem-solving, and communication skills needed to use artificial intelligence responsibly in academic, professional, and everyday contexts.

Credits - One Semester (0.5 Carnegie unit / CA: 5 credits) | Prerequisites: None

Course Outline

Unit 1: What Is AI?

- 1.1 A Brief History of "Thinking Machines"
- 1.2 The Turing Test
- 1.3 Early Successes (and Failures)
- 1.4 AI vs. People
- 1.5 What Does an AI Model Actually Know?

Unit 2: How LLMs Work

- 2.1 What Is a Model?
- 2.2 How Can LLMs Do Things Never Seen Before?
- 2.3 Hallucinations, Errors, and How to Catch Them
- 2.4 What LLMs Are Good At
- 2.5 The AI Family Tree
- 2.6 How Machines Learn — A Closer Look at Machine Learning

Unit Objectives

In Unit 1 students will:

Explore the history of artificial intelligence, from early ideas about thinking machines to the development of modern AI systems.
Examine the Turing Test and evaluate its role in measuring machine intelligence and its limitations.
Investigate early successes, challenges, and periods of rapid growth that shaped the development of artificial intelligence.
Compare how humans and AI process information, highlighting the differences between pattern recognition and human understanding.
Analyze what it means for an AI model to "know" information and how AI systems generate responses.
Evaluate the strengths and limitations of AI knowledge, including why AI can produce accurate information in some situations and unreliable information in others.

In Unit 2 students will:

Explore what AI models are, how they are trained, and how they differ from traditional computer programs.
Examine how large language models generate original responses without understanding information the way humans do.
Investigate why AI systems produce errors and hallucinations and apply strategies for identifying unreliable outputs.
Evaluate the strengths of large language models for tasks such as writing, research, coding, analysis, and problem-solving.
Compare different branches of artificial intelligence, including machine learning, deep learning, and generative AI.
Understand how machine learning systems learn from data and how the training process influences AI performance and limitations.

Course Outline

Unit 3: Ethics, Risk, and Trust

- 3.1 Can We Trust AI?
- 3.2 Bias and Fairness
- 3.3 Before We Go Further: A Brief Word on Cheating
- 3.4 Deepfakes and Misinformation
- 3.5 AI and Copyright
- 3.6 Privacy and Consent

Unit 4: Becoming an AI Power User

- 4.1 AI Skill #1: Prompt Engineering
- 4.2 Advanced Prompting: What Changed and What Remains Useful
- 4.3 Lexical vs. Semantic Search
- 4.4 Retrieval Augmented Generation (RAG)

Unit 5: Beyond Chat

- 5.1 Image Models
- 5.2 Audio Models
- 5.3 Video, Film, and the Lifelike Avatar
- 5.4 If It's AI, Is It Art?
- 5.5 Advanced Tools and Agents
- 5.6 You Can Program, Even If You Don't Know How

Unit 6: Your Future with AI

- 6.1 Other Types of AI
- 6.2 The Future of Work
- 6.3 The Experience Gap
- 6.4 Proceeding With Caution

Unit Objectives

In Unit 3 students will:

Evaluate when AI systems can be trusted and apply strategies for using AI responsibly in different situations.

Examine how bias can develop in AI systems and analyze its impact on fairness and decision-making.

Understand the role of academic integrity when using AI tools and distinguish between appropriate and inappropriate uses.

Investigate how AI-generated misinformation and deepfakes are created, shared, and evaluated for credibility.

Explore copyright, ownership, and legal considerations related to AI-generated content.

Examine how AI systems collect and use personal data and identify strategies for protecting privacy and obtaining consent.

In Unit 4 students will:

Develop effective prompt engineering skills by writing clear instructions, setting constraints, assigning roles, and refining prompts.

Evaluate advanced prompting strategies and determine which techniques improve AI-generated responses.

Compare keyword-based and meaning-based search methods and explain how they influence the information AI systems retrieve.

Explore retrieval-augmented generation (RAG) and understand how AI systems use external documents and data to generate more accurate responses.

In Unit 5 students will:

Explore how AI image models generate visual content and evaluate the strengths and limitations of AI-generated images.

Examine how AI systems create speech, voices, and music while identifying techniques for recognizing synthetic audio.

Investigate AI-generated video and digital avatars and evaluate their impact on communication, media, and the reliability of visual evidence.

Analyze questions of authorship, originality, and creativity in AI-generated art and media.

Explore AI agents that can complete tasks on a user's behalf and evaluate their capabilities and limitations.

Apply AI tools to design and build a simple interactive application while learning how AI can support the software development process.

In Unit 6 students will:

Explore emerging areas of artificial intelligence, including robotics, reinforcement learning, diffusion models, and artificial general intelligence (AGI).

Examine how AI is transforming the workplace and evaluate the opportunities and challenges created by automation.

Analyze how AI is changing the development of expertise and the skills needed for future careers.

Evaluate AI safety, alignment, and other long-term challenges while considering the responsible development and use of artificial intelligence.