NAAIC

Resource provided by

MARICOPA COMMUNITY COLLEGES

Course Outline Details

Introduction to Artificial Intelligence

This introductory course provides hands-on experience with artificial intelligence concepts, techniques, and real-world applications.

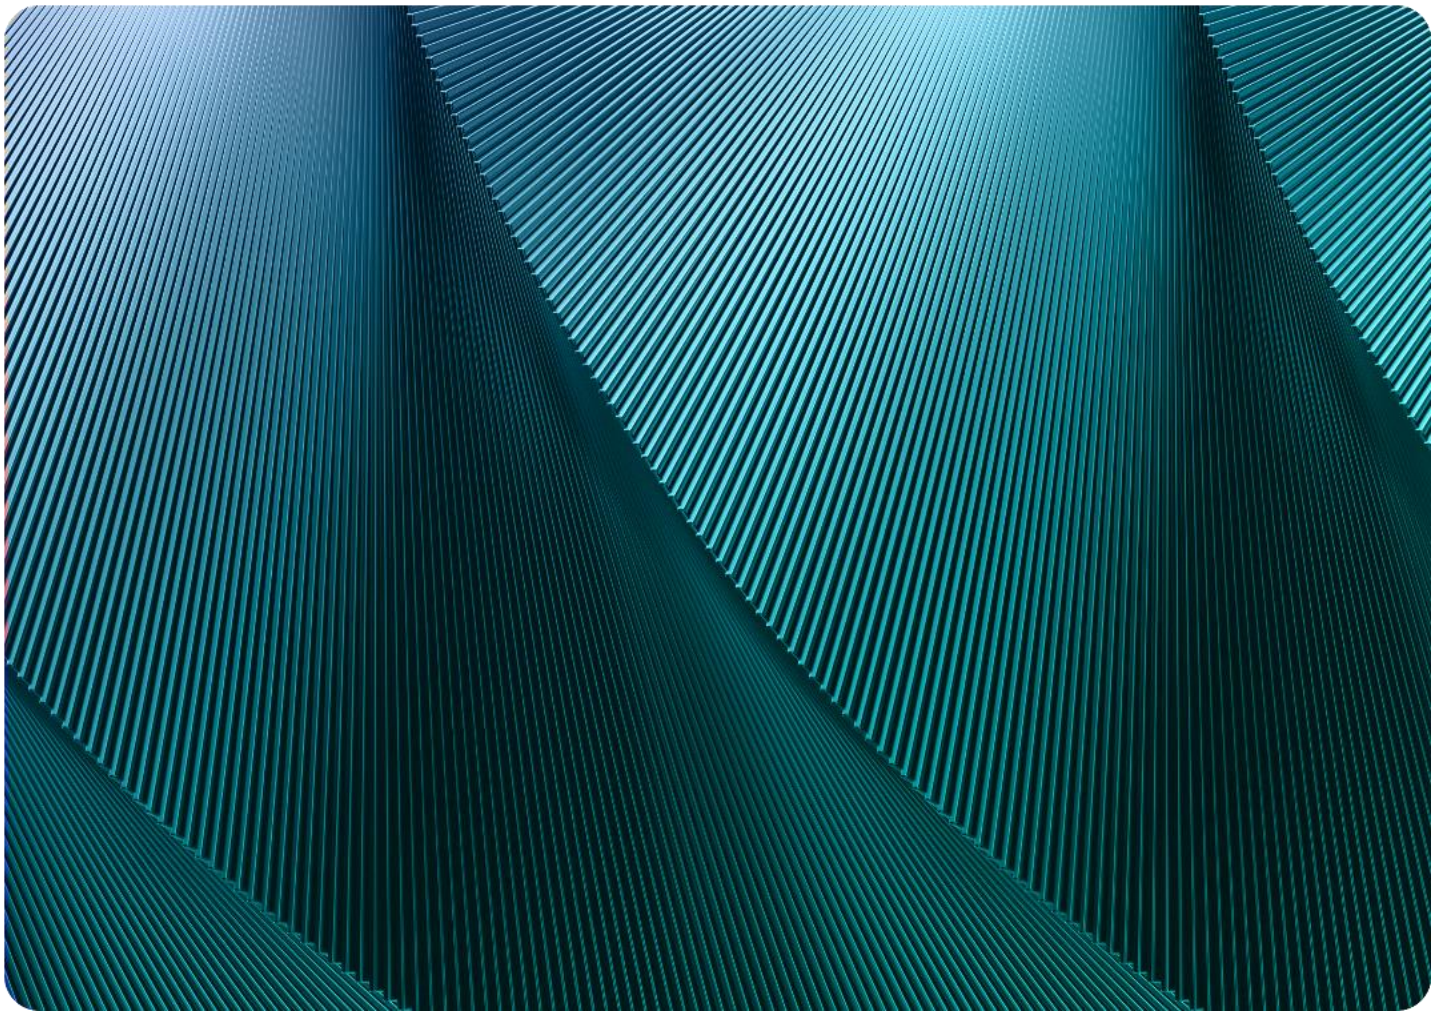

NAAIC

Course Outline

Week starting on / Module	Module Topic	Assignments
Week 1	Lecture 0 - Introduction/Syllabus Lecture 1 - Introduction to AI	1. Discussion 0 - Intro and Goal Setting Due 2. Syllabus Quiz Due 3. Lab 1 Assigned
Week 2	Lecture 2 - Intro to Data Lecture 3 - Intro to Computer Vision Lecture 4 - Intro to Natural Language Processing	1. Quiz 1 - Data Due 2. Quiz 2 - Computer Vision Due 3. Quiz 3 - Natural Language Processing Due
Week 3	Lecture 5 - Machine Learning Modeling and Workflow	1. Lab 2 Due 2. Quiz 4 - ML Modeling and Workflow Due
Week 4	Lecture 6 - Emerging Technologies	1. Emerging Tech Presentation Due
Week 5	Lecture 7 - Large Language Models and Text-to-Image Models	1. Lab 3 Due
Week 6	Lecture 8 - AI Industry Guest Speaker and the Future of Success in AI	1. Discussion 1 - AI Jobs in Industry Due 2. Lab 4 Due
Week 7	Lecture 9 - AI Ethics	1. AI Ethics Presentation Due 2. Lab 5 Due
Week 8	Final Exam	1. Final Exam Due

NAAIC

Course Outline Details

Lecture 1

Back to Outline

Introduction to AI, Course Orientation, and Personal Goal Setting

By the end of this lecture, students should be able to:

- Understand the course structure, learning expectations, grading policies, and available resources.
- Define Artificial Intelligence (AI), explain how it differs from traditional programming, and identify its basic working principles, such as learning from numerical data.
- Recognize key AI subfields, including Machine Learning and Deep Learning, and differentiate between their main types (Supervised, Unsupervised, and Reinforcement Learning).
- Articulate their personal motivations for studying AI and establish actionable, semester-long goals with strategies for overcoming potential obstacles.

Key Activities & Pedagogical Approaches:

- **Syllabus Review:** Students will review the course syllabus to understand expectations, a document presented as a "contract" between instructor and student.
- **Interactive Lecture:** Engagement with an introductory lecture on AI, covering its definition, historical context, and differentiation from standard software. This includes an activity where students identify AI versus non-AI examples.
- **Goal Setting Exercise:** Students will read an article on goal setting and reflect on their personal and academic objectives for the course and semester, including identifying potential obstacles and mitigation strategies.
- **Online Discussion:** Participation in a discussion board to introduce themselves, share their reasons for taking the course, and discuss their approaches to achieving their defined goals (focusing on actions rather than the goals themselves).

Flow & Connection to Previous/Next Lecture:

This is the inaugural week, setting the stage for the entire 8-week course. It provides the essential vocabulary, core concepts of AI, and course expectations necessary for students to engage with more specialized topics and hands-on lab work in subsequent weeks. The goal-setting activity also aims to establish a proactive learning mindset for the semester.

Lecture 2

Back to Outline

The Fundamental Role and Characteristics of Data in Artificial Intelligence

By the end of this lecture, students should be able to:

- Define "data" and articulate its critical importance as the foundational element for AI systems.
- Differentiate between structured data (organizable in rows and columns, e.g., databases, spreadsheets) and unstructured data (e.g., images, audio, text files, emails), recognizing that unstructured data forms the majority of enterprise data.
- Distinguish between qualitative data (descriptive, categorical, e.g., labels, names) and quantitative data (numerical, measurable, e.g., sensor readings, counts).
- Understand and identify key data characteristics such as features, association (the meaningful connection between data points), data collection frequency, relevant time period for data, and observable trends, and explain their relevance to building machine learning models.
- Appreciate the value of data exploration and visualization as methods to comprehend data patterns and communicate insights effectively.

Key Activities & Pedagogical Approaches:

- **Interactive Lecture:** Students will engage with a presentation covering the definition of data, its various types (structured/unstructured, qualitative/quantitative), essential characteristics (features, association, frequency, period, trend), and the significance of data exploration and visualization. The lecture includes examples like weather data related to health cases, car attributes for price prediction, and user preferences for movie recommendations.

- **Problem-Solving Scenarios ("Quest Alerts"):** The lecture incorporates hypothetical AI challenges (e.g., optimizing farming, preventing elephant poaching, preserving historical structures) to encourage students to brainstorm relevant data features.
- **Conceptual Reinforcement:** An example of an AI playing Rock, Paper, Scissors is used to illustrate how AI systems learn and improve performance with more data (more gameplays).
- **Assessment Quiz:** Students will complete "Quiz 1 - Data," which presents scenarios (e.g., monitoring Arctic Sea ice health for NASA, predicting manufacturing tool maintenance) requiring them to identify appropriate data types, relevant features, assess data frequency issues, and interpret data trends.

Flow & Connection to Previous/Next Lecture:

This module logically follows a general introduction to AI by delving into the critical component of data. A solid understanding of data types, characteristics, and handling is essential before students can effectively learn about specific AI domains like Computer Vision and Natural Language Processing (covered later this week or in subsequent weeks).

Back to Outline

Lecture 3

Fundamentals of Computer Vision, Image Data Representation, and Core CV Tasks

By the end of this lecture, students should be able to:

- Define Computer Vision (CV) as the AI field enabling computers to interpret and understand information from images and videos.
- Explain that digital images are composed of pixels and describe how pixel values represent intensity in grayscale images (0 for black, 255 for white) and color in RGB images (three values from 0-255 for Red, Green, Blue).
- Identify and differentiate key CV tasks such as Image Classification (assigning a label to an image), Classification + Localization (assigning a label and drawing a bounding box around a single object), Object Detection (identifying and locating multiple objects with bounding boxes), and Instance Segmentation (identifying and delineating each object instance at the pixel level).

- Distinguish between Statistical (Traditional) Computer Vision (which relies on manually defined image features like edges, corners, or Haar-like features) and Deep Learning based Computer Vision (which automates the feature detection and extraction process).
- Recognize critical factors for acquiring high-quality input image data for CV, including photo resolution, lighting conditions, object angles and distance, image size consistency, and managing variations in viewpoint, scale, and illumination.

Key Activities & Pedagogical Approaches:

- **Illustrated Lecture:** Students will learn through a presentation that visually explains the core concepts of Computer Vision, how computers "see" images as numerical data, common CV tasks with illustrative examples, and the differences between traditional and modern deep learning approaches.
- **Interactive Examples:** The lecture poses questions like "What do you see?" versus "What do computers see?" to highlight the complexities CV aims to solve. It also includes examples of how different colors are represented by RGB values.
- **Assessment Quiz:** "Quiz 2 - Computer Vision" will test students' knowledge of pixel values, the basic building blocks of digital images, identifying visual features (e.g., edges, corners, areas of interest from images), the advantages of deep learning in CV, and the suitability of datasets for specific CV tasks based on image characteristics.

Flow & Connection to Previous/Next Lecture:

This module introduces a significant AI specialization. It directly builds upon the foundational understanding of "Data" (Lecture 2), as Computer Vision is entirely reliant on processing image data. The concepts learned here (pixels, features, classification) will be essential for any subsequent practical labs in Computer Vision (such as the "Hands-On Look at Computer Vision" Lab 1) and for understanding more advanced CV architectures.

NAAIC

Back to Outline

Lecture 4

Core Concepts of Natural Language Processing (NLP), Common Tasks, Challenges, and Basic Processing Pipeline

By the end of this lecture, students should be able to:

- Define Natural Language Processing (NLP) as a subfield of AI focused on the interaction between computers and human (natural) languages.
- Identify a variety of common NLP tasks (e.g., sentiment analysis, text classification, chatbots/virtual assistants, text extraction, machine translation, text summarization, auto-correct, speech recognition) and their real-world applications (e.g., spam filtering, medical diagnosis from speech).
- Recognize significant challenges in NLP, including linguistic ambiguity (syntax vs. semantics), variations in grammar and dialects, the nuances of spoken language (accents, errors), and difficulties in interpreting sarcasm or figurative language.
- Describe the fundamental steps in a traditional (statistical) NLP pipeline, such as sentence segmentation, word tokenization, part-of-speech (POS) tagging, lemmatization, stop word removal, dependency parsing, and Named Entity Recognition (NER).
- Understand the difference between statistical NLP (feature-based) and Deep Learning NLP (which aims to automate feature creation for a more end-to-end approach).
- List common sources for acquiring data for NLP training, such as publicly available datasets (e.g., Twitter feeds, Yelp reviews), web scraping, and social media platforms.

Key Activities & Pedagogical Approaches:

- **Comprehensive Lecture:** Students will be introduced to NLP through a presentation defining the field, outlining its tasks and applications, discussing its inherent challenges, detailing the steps of a traditional processing pipeline using a consistent example, and contrasting statistical versus deep learning methods. The lecture also covers data sources for NLP.

- **Example-Driven Explanations:** The lecture uses concrete examples for NLP tasks (e.g., sentiment analysis of product reviews, spam email classification, chatbot interaction) and pipeline steps (processing a paragraph about London).
- **Assessment Quiz:** "Quiz 3 - Natural Language Processing" will evaluate students' comprehension of NLP challenges (spoken and written), common uses of NLP, methods for text segmentation (by sentence, by word), identifying parts of speech, recognizing stop words, understanding Named Entity Recognition (NER), and common data sources for NLP training.

Flow & Connection to Previous/Next Lecture:

This lecture logically follows a general introduction to AI by delving into the critical component of data. A solid understanding of data types, characteristics, and handling is essential before students can effectively learn about specific AI domains like Computer Vision and Natural Language Processing.

Lecture 5

Back to Outline

Understanding and practically implementing a comprehensive Machine Learning (ML) Workflow. The core stages covered are Problem Scoping, Data Acquisition, Data Exploration, Modeling, Evaluation, and Deployment. Students will gain hands-on experience using Microsoft Azure Machine Learning Studio to build, train, and evaluate a predictive model

By the end of this lecture, students should be able to:

- Clearly distinguish between an ML **algorithm** (the procedure or set of rules) and an ML **model** (the output artifact generated by running an algorithm on data).
- Describe each key stage of a standard ML workflow and its importance:
    - **Problem Scoping:** Defining the problem, identifying the target audience, establishing boundaries (resources, requirements), and assessing potential risks and ethical considerations (e.g., bias, need for human judgment).
    - **Data Acquisition:** Identifying and obtaining relevant data from various sources such as online repositories (Google Dataset Search, Kaggle), web scraping, or manual collection.
    - **Data Exploration:** Analyzing datasets to understand their main characteristics, often employing data visualization techniques to discover patterns, spot anomalies, and test hypotheses (Exploratory Data Analysis - EDA).

  - **Modeling:** Selecting appropriate ML algorithms based on data type (quantitative/numerical vs. qualitative/categorical), preparing data through techniques like one-hot encoding (for converting categorical features to numerical) or binning (numerical to categorical), and training the model.
  - **Evaluation:** Assessing model performance using suitable evaluation metrics (e.g., Mean Absolute Error for regression tasks) by comparing predicted outcomes against actual data.
  - **Deployment:** Understanding the activities that make an ML system available for use, and considering factors like computational resources, hardware limitations, latency, privacy, and the need for ongoing maintenance.
- Gain practical experience using Microsoft Azure Machine Learning Studio to create an end-to-end ML pipeline, including importing data, preprocessing it, training a model (specifically Linear Regression for automobile price prediction), and evaluating its performance.
- Independently apply the learned workflow by selecting a new dataset (from provided options like Insurance, Cancer, or Real Estate data), choosing relevant features, cleaning the data, selecting an appropriate regression algorithm, training a model, and evaluating its outcomes within the Azure ML Studio environment.

Key Activities & Pedagogical Approaches:

- **Conceptual Lecture ("ML Modeling and Workflow"):** Students will engage with a presentation that systematically breaks down the ML workflow. This includes differentiating models from algorithms, discussing when *not* to use ML, outlining rules for problem scoping, detailing various data acquisition methods, explaining data exploration using EDA, covering data type considerations for model selection (e.g., Decision Trees for categorical, Linear Regression for numerical), data transformations (one-hot encoding, binning), model evaluation strategies, and deployment considerations.

- **Hands-on Lab ("Implementing a ML Workflow" using Microsoft Azure):**
    - **Guided Project:** Students will first work through a step-by-step tutorial to construct a Linear Regression model in Azure Machine Learning Studio. The goal is to predict automobile prices using a sample dataset. This part of the lab involves setting up an Azure for Students account, learning the Azure ML Studio interface, building a pipeline by dragging and dropping components, importing and preparing data, training the linear regression model, and evaluating it, focusing on the Mean Absolute Error (MAE).
    - **Independent Project:** Following the guided exercise, students will choose one of three provided datasets (Insurance, Cancer, or Real Estate). They will then create their own complete ML workflow in Azure, which includes uploading their chosen dataset, selecting relevant columns for analysis, cleaning missing data (by removing rows), splitting the data for training and testing, selecting a regression algorithm from a list (e.g., Linear Regression, Poisson Regression, Decision Forest Regression, Neural Network Regression), training the model to predict a specified target variable (e.g., "charges" for Insurance, "TARGET_deathRate" for Cancer), and evaluating the model's performance.
- **Video-Assisted Learning:** The lab guide provides links to YouTube tutorials for key steps like Azure student account sign-up, general Azure ML Pipeline usage, and adding custom datasets to Azure ML Studio.
- **Reflective Learning:** Students are prompted to document their understanding (e.g., summarizing Linear Regression), record specific model outputs (predicted vs. real prices, MAE), and answer questions regarding their choices and findings in a lab template.

Flow & Connection to Previous/Next Lecture:

- This lecture is pivotal as it synthesizes many previously discussed concepts (AI fundamentals, data types, the importance of data) into a cohesive, practical framework for building machine learning solutions. The hands-on experience with Azure ML Studio bridges theory and application, demystifying the process of creating and evaluating ML models. This experience is foundational for any further exploration into more advanced machine learning techniques, specialized AI applications, or capstone projects.

NAAIC

Lecture 6

Back to Outline

An overview of key Emerging Technologies that are integral to or supportive of Artificial Intelligence. Covered topics include Computer Hardware (CPUs, GPUs, specialized AI chips), Autonomous Vehicles, Quantum Computing, the Internet of Things (IoT), and 5G Networks. The week culminates in an "Emerging Tech Presentation" assignment where students research and present on an AI-related emerging technology of their choice

By the end of this lecture, students should be able to:

- Identify and describe the fundamental concepts, significance, and AI-relevance of several emerging technologies:
    - **Computer Hardware:** Understand the roles of CPUs, GPUs, and Integrated Circuits (ICs), the concept of Moore's Law, and the emergence of specialized AI hardware like TPUs and FPGAs.
    - **Autonomous Vehicles:** Describe the necessary sensor technologies (LiDAR, radar, cameras), the SAE levels of automation (0-5), and discuss major players, benefits, and challenges.
    - **Quantum Computing:** Grasp basic principles like qubits, and the potential of quantum algorithms (e.g., Shor's, Grover's) to solve complex problems, including their application to AI challenges like dimensionality reduction.
    - **Internet of Things (IoT):** Explain the core components (sensors, connectivity, data processing, UI), common communication protocols, and the impact of IoT on data generation and AI applications.
    - **5G Networks:** Recognize key features like high speed and low latency, and understand 5G's role in enabling advanced AI applications, edge computing, and massive IoT connectivity.
- Conduct independent research on an AI-related emerging technology that interests them, analyzing it from multiple perspectives including advantages and potential issues.
- Develop and deliver a concise (4-5 minutes) recorded presentation using screen sharing and webcam, clearly explaining their chosen topic, their standpoint, and what they learned.

- Effectively engage in peer review by providing constructive, concept-building feedback on classmates' presentations via a discussion board.

Key Activities & Pedagogical Approaches:

- **Lecture on Emerging Technologies:** Students will learn about the aforementioned technologies through a dedicated lecture presentation that details their workings, key players (e.g., Intel/AMD for CPUs, NVIDIA for GPUs, Waymo/Tesla for autonomous vehicles), current challenges, benefits, and specific connections to AI.
- **"Emerging Tech Presentation" Assignment:**
    - **Topic Selection & Research:** Students choose an AI-related emerging technology, which can be an existing technology undergoing new developments or an imaginative prototype. A list of suggested online resources (e.g., Medium, TechCrunch, Reddit) is provided for idea generation.
    - **Presentation Development:** Students create a visual presentation (e.g., PowerPoint), with a hint from the instructor to favor bullet points over large blocks of text for clarity.
    - **Recording & Sharing:** Presentations must be recorded using **Loom**, showcasing both the presentation material (full screen) and the student's webcam. The recording is limited to 4-5 minutes, aligning with Loom's free version constraints. The shareable Loom embed code is then posted to a designated Canvas discussion board.
- **Peer Review:** Students are required to watch and respond to at least two other students' presentations, with specific instructions to build on concepts rather than simply agreeing or disagreeing.

Flow & Connection to Previous/Next Lecture:

- This lecture broadens the students' understanding of the technological landscape that AI operates within and depends upon. The presentation assignment serves as a practical application of research and communication skills, allowing students to delve deeper into an area of personal interest related to AI. The project can act as a bridge to future discussions on specialized AI applications, ethical considerations of new technologies, or career paths in the AI field.

NAAIC

Lecture 7

Back to Outline

Introduction to Generative AI, focusing on Prompt Engineering techniques for Large Language Models (LLMs) with an emphasis on ChatGPT, and for Text-to-Image models, primarily Stable Diffusion

By the end of this lecture, students should be able to:

- Define Generative AI (GenAI) and describe its capability to create a wide variety of data, including images, video, audio, and text.
- Explain what Language Models and Large Language Models (LLMs) are, understand their probabilistic nature in generating word sequences, and list common applications such as speech recognition, machine translation, and chatbots.
- Understand and apply fundamental Prompt Engineering principles for LLMs, including the core elements of a prompt (instruction, context, input data, output indicator), and differentiate between techniques like zero-shot and few-shot prompting.
- Practice crafting effective prompts for an LLM (ChatGPT) for both descriptive and creative tasks, iteratively refine prompts based on model responses, and apply an advanced prompting technique such as Chain-of-Thought.
- Define Text-to-Image models as AI systems that generate visual content from textual descriptions and list their applications, such as creating illustrations or assisting in design.
- Identify and utilize key components of an effective prompt for text-to-image generation, including subject, medium, style, artist, resolution, additional details, color, and lighting.
- Gain hands-on experience generating images using a text-to-image model (Stable Diffusion) by constructing detailed positive prompts, employing negative prompts to refine outputs, and creating images based on both provided scenarios and original ideas.

Key Activities & Pedagogical Approaches:

- **Interactive Lectures:** Students will attend two lectures: one on "Large Language Models" and another on "Text-to-Image Models". These lectures cover the definition, history, common examples, applications, and specific prompt engineering techniques relevant to each type of model.

- **Hands-on Lab (Lab 3 - "Generating with ChatGPT and StableDiffusion"):**
    - **LLM Interaction (ChatGPT):** Students will create an OpenAI account to access ChatGPT. They will practice writing descriptive prompts on a topic of interest, crafting creative prompts for imaginative scenarios, and iteratively fine-tuning prompts for a less familiar topic to achieve deeper understanding. A specific exercise involves understanding and applying Chain-of-Thought prompting to a given problem.
    - **Text-to-Image Generation (Stable Diffusion):** Students will use a web-based Stable Diffusion platform (or an alternative like Craiyon) to generate images from text prompts. They will learn about and apply negative prompting by comparing images generated with and without negative prompts. The lab includes structured "Art Gallery Submissions" where students generate images for specific themes (e.g., outer space, their favorite season) and one freestyle creation, focusing on using multiple prompt anatomy elements.
- **Peer Sharing and Community Building:** Students will post their generated images from the lab to a Canvas discussion board titled "AIM Art Gallery", allowing them to share their work and see the creations of their peers.
- **External Resource Engagement:** The curriculum encourages students to consult external resources like promptingguide.ai and stable-diffusion-art.com for deeper insights into prompt engineering techniques.

Flow & Connection to Previous/Next Lecture:

- This lecture's focus is on Generative AI, particularly LLMs and Text-to-Image models, builds directly on the foundational AI concepts, data literacy (Lecture 2), and potentially introductory NLP concepts (Lecture 4) from previous lectures. The hands-on experience with prompt engineering is a critical skill in interacting with modern AI systems. This lecture provides practical skills and understanding that can be applied to more advanced AI topics, projects, or discussions on the broader impacts and ethics of generative technologies in subsequent lectures.

NAAIC

Back to Outline

Lecture 8

Professional landscape of Artificial Intelligence, exploring AI job roles, current industry trends, and essential career development strategies. It also emphasizes the practical application of AI tools, specifically ChatGPT, for enhancing productivity and managing challenging daily tasks. Insights from an industry professional are integrated into career exploration activities

By the end of this lecture, students should be able to:

- Analyze the current state of the AI industry, including job market trends (e.g., growth in AI job postings, sector-specific demand), private investment focus areas, and organizational adoption rates, by interpreting data from sources like the Stanford HAI AI Index Report.
- Identify and differentiate between key career paths in AI, such as Data Analyst (focusing on data cleaning, exploratory data analysis, visualization, labeling, and curating data), Machine Learning Engineer (designing, training, deploying, and monitoring ML models), and Data Scientist (hypothesis testing, developing predictive models, conducting experiments), including their typical educational requirements.
- Articulate the general public's responsibilities in an AI-driven world, encompassing basic AI understanding, data literacy, ethical awareness, and familiarity with AI tools.
- Develop strategies for gaining relevant experience for an AI career, such as working on personal/school projects, participating in hackathons and coding challenges, contributing to open-source projects, and seeking internships.
- Conduct thorough research on a potential future career (AI-related or how AI might impact another field), detailing educational and experience prerequisites, typical job responsibilities, essential skills and technologies, and prominent industries or employers, integrating insights from a guest lecture.
- Utilize ChatGPT effectively to address common challenging personal and professional situations, including drafting difficult emails with appropriate tone, summarizing complex or technical information concisely, and organizing multi-step tasks or projects for better clarity and prioritization.
- Practice iterative prompt refinement with ChatGPT to achieve more precise and useful outputs for specific tasks.

Key Activities & Pedagogical Approaches:

- **Industry Insights Lecture ("AI Jobs in Industry"):** Students will be presented with current data and trends in the AI job market, drawing from resources like the Stanford HAI AI Index Report 2024, Lightcast, and McKinsey surveys. The lecture will cover AI job posting percentages by geographic area and U.S. sector, private investment in AI by focus area, U.S. AI labor demand by state, median salaries for various developer roles, and rates of AI adoption by organizations. It will also detail different AI career paths, the importance of building a portfolio, networking, and how to analyze job requisitions.
- **Guest Lecture Integration:** Insights from a provided YouTube video featuring an industry professional will serve as a guest lecture. Students will be expected to incorporate learnings from this lecture into their career research and discussion.
- **Career Research and Discussion Assignment ("AI Jobs in Industry Discussion"):** Students will conduct research into a career path of interest (AI-focused or otherwise, including how AI might reshape existing roles). Their research should cover:
    - Educational and Experience Requirements.
    - Job Responsibilities.
    - Skills and Technologies to master.
    - Key Industries and Employers, including growth areas.
    - A specific reflection on the most interesting insight gained from the guest speaker and its influence on their career perspective. Students will post their findings (minimum 3 sentences per topic) to a discussion board and engage constructively with at least two classmates' posts. Guidance on research resources like university websites, job listings, professional associations, and tech blogs is provided.

- **Practical AI Lab ("Overcoming Challenging Tasks with ChatGPT"):** This lab focuses on using ChatGPT as a tool to manage difficult real-world tasks. Students will:
    - Craft a detailed prompt for ChatGPT to help write a challenging email (e.g., resolving conflict, asking for a raise), focusing on specifics like recipient, main points, desired tone, and outcome.
    - Prompt ChatGPT to summarize a piece of complex or technical information for a non-technical audience, specifying the desired length and focus.
    - Use ChatGPT to organize a chaotic, multi-step project or to-do list, detailing tasks, deadlines, and dependencies to get help with prioritization and planning.
    - Refine one of their previous prompts for greater specificity and compare the original ChatGPT response with the refined one.
    - Define a personal challenging task, explain its impact, and reflect on how ChatGPT could help.
- Write a brief reflection on how ChatGPT helped in tackling these tasks, noting any reduction in stress or time saved.

Flow & Connection to Previous/Next Lecture:

- This lecture provides a crucial bridge between foundational AI knowledge and its real-world implications for careers and practical daily use. The focus on job market realities, career development, and the utility of AI tools like ChatGPT helps students contextualize their learning and consider future paths. This lecture prepares students for more advanced career planning, ethical discussions, or specialized AI electives.

NAAIC

Back to Outline

Lecture 9

A comprehensive exploration of key Ethical Implications in Artificial Intelligence and Robotics. This module involves a lecture on prevalent ethical debates and a significant student-led research presentation on a chosen AI ethics topic

By the end of this lecture, students should be able to:

- Identify and articulate the nuances of major ethical debates in AI, such as Privacy and Surveillance, Manipulation of Behavior, Opacity of AI Systems (the "black box" problem), Bias in Decision Systems, Human-Robot Interaction, the impact of Automation on Employment, the concept of Machine Ethics (e.g., ethics for machines themselves), the notion of Artificial Moral Agents, and the potential societal disruption of Singularity/Superintelligence.
- Analyze various ethical frameworks and guidelines for AI being developed and discussed globally.
- Conduct in-depth research into a specific AI ethics topic selected from a provided list (based on the Stanford Encyclopedia of Philosophy's entry on AI Ethics).
- Develop and deliver a concise (4-5 minutes), well-structured recorded presentation that:
    - Provides a clear overview of the chosen ethical topic, understandable to a non-expert, including current developments.
    - Discusses various stances on the topic, including potential advantages, disadvantages, and the audiences likely to be benefited or harmed.
    - Articulates and justifies their personal standpoint on the ethical issue.
    - Summarizes what they learned from studying the topic and lists their research sources.
- Engage in thoughtful peer review by providing constructive, recorded **Loom** feedback (minimum 2 minutes long) on the presentations of at least two classmates, focusing on building upon the presented concepts.

Key Activities & Pedagogical Approaches:

- **AI Ethics Lecture:** Students will attend a lecture covering common societal fears surrounding AI (e.g., bias, opacity, job loss), detailing the main ethical debates within the field (drawing from the Stanford Encyclopedia of Philosophy), examining global AI ethics guidelines (including common principles like transparency, justice, and non-maleficence), and presenting a critical view on the feasibility of universally "Ethical AI". An in-class activity might involve students analyzing ethical policies of different countries.

- **"AI Ethics Presentation" Assignment:**
    - **Topic Selection & Research:** Students choose one ethical topic related to AI from a list (e.g., Privacy, Bias, Automation, Singularity) and conduct detailed research using provided resources like the Stanford Encyclopedia of Philosophy and links to AI ethics organizations (e.g., Turing Institute, AI Now).
    - **Presentation Creation & Recording:** Students develop a visual presentation (e.g., PowerPoint, with a suggestion to use bullet points for clarity) and record themselves presenting it using **Loom**. The recording must include both the shared screen (in full-screen presentation mode) and the student's webcam and be between 4 and 5 minutes long.
- **Submission & Peer Feedback:** The Loom presentation is submitted by embedding it into a class discussion board. Students then provide video-recorded Loom feedback (minimum 2 minutes) to two peers, aiming to expand on ideas rather than merely agreeing or disagreeing.
- **Guided Audio Model Creation (Sound Events):** Students will follow a provided video tutorial to build their first audio classifier in Teachable Machine. This model will be trained to distinguish background noise from specific sounds like snaps, claps, and whistles (or alternative sounds like tongue clicks if a student cannot whistle). Students will then test this model's robustness by varying input conditions, such as combining sounds, changing microphone distance, or altering background noise, and recording their observations.
- **Independent Audio Model Creation (Language Classification):** Students will embark on creating a more complex audio classifier designed to identify spoken language. They will:
    - Create a new audio project in Teachable Machine.
    - Record samples of background noise.
    - For three distinct languages of their choice, record themselves speaking a list of eight specified phrases (e.g., "Hello," "Thank you"). Google Translate is suggested as a tool to hear correct pronunciations.
    - Train the model and then test its ability to classify new words (e.g., "Please," "Welcome") spoken in each of the chosen languages, adjusting parameters like the "overlap factor" during testing.

- **Conceptual Learning and Reflection:** Students will read an excerpt from an external article about speech-to-text models to gain insight into audio signal processing. They will then write a summary of what they learned about how microphones capture sound as wavelengths and how ML models use these digital signals.

Flow & Connection to Previous/Next Lecture:

- Positioned towards the end of the course, this lecture on AI ethics encourages students to critically reflect on the societal and moral dimensions of the AI technologies and applications studied in previous lectures. The presentation assignment is a significant capstone-like activity, allowing for deep engagement with a chosen ethical concern; it emphasizes responsible innovation and is crucial for a holistic understanding of AI's impact.

Final Exam

Back to Outline

Comprehensive summative assessment of core Artificial Intelligence concepts, principles, applications, and ethical considerations covered throughout the course

Upon completion of the course and this exam, students should be able to demonstrate their understanding and ability to:

- Identify and describe the major steps involved in a standard Machine Learning workflow.
- Explain the synergistic relationship between advancements in foundational technologies, such as Integrated Circuits (ICs) and 5G Networks, and the capabilities of Artificial Intelligence systems.
- Recall and apply basic concepts from emerging technology areas relevant to AI, including quantum computing (distinguishing qubits from bits) and autonomous vehicles (identifying sensor technologies and levels of driving automation).
- Recognize the key components of an effective prompt for text-to-image generative AI models.
- Differentiate between structured and unstructured data, applying this understanding to common examples like barcodes and social media content.
- Identify typical tasks and applications handled by Large Language Models (LLMs), such as question answering and text summarization.

- Analyze simplified, scenario-based problems related to Computer Vision (CV) and Natural Language Processing (NLP), including identifying appropriate AI applications, factors influencing data quality for CV, and basic feature identification for CV tasks.
- Apply basic NLP text processing concepts (e.g., sentence segmentation) and critically evaluate the nuances and potential limitations of AI systems, such as in sentiment analysis tasks.
- Identify relevant data analytics that would be needed to inform decision-making in a given scenario.
- Articulate and explain significant ethical dilemmas associated with AI, particularly concerning user privacy and bias in AI-driven decision-making systems.

Key Activities & Pedagogical Approaches:

- **Final Examination:** A comprehensive, timed (170 minutes) final exam comprising 20 questions designed to evaluate students' understanding of the course material. The exam is administered and submitted through a learning management system.
- **Assessment Methods:** The exam utilizes a mix of question formats to assess both knowledge recall and application, including:
    - Multiple-select questions.
    - Short answer/explanation questions requiring 2-3 sentences.
    - Multiple-choice questions.
- True/False questions. A significant portion of the exam involves scenario-based questions where students apply AI concepts to a hypothetical situation, such as designing an AI-powered university parking system or analyzing an NLP-based review system.

Flow & Connection to Previous/Next Lecture:

- The final exam serves as the culminating summative assessment for the course. It is designed to measure the overall achievement of course learning objectives based on content delivered throughout all preceding lectures. It marks the conclusion of the formal instruction for the course