07-280 Syllabus

Overview

Course Title: Artificial Intelligence and Machine Learning I

Course Number: 07-280

Course Website: https://www.cs.cmu.edu/~07280

Instructors: MLD + CSD

S26: Pat Virtue, Nihar Shah

Units: 12

Frequency: Every semester, starting S26

Prerequisites: (15122) and (21240 or 21241 or 18202 or 21254) and (15151 or 21127 or 21128) Minimum grade C in each

Select prereq course names:

15-122 Principles of Imperative Computation

21-241 Matrices and Linear Transformations

21-127 Concepts of Mathematics or 15-151 Mathematical Foundations of Computer Science

Corequisites: 21-122 Integration and Approximation

36-225 Introduction to Probability Theory (or 36217 or 15259 or 21325 or 15359 or 36218 or 36219 or 36235)

Antirequisite: 15281, 10301, 10601, 10315, 18461, 18661

Anti-req Prohibits: 15181, 15281, 10301, 10601, 10315, 18461, 18661

Textbook: No required textbook

Description: This course provides an integrated introduction to artificial intelligence and machine learning that bridges core AI methods with modern approaches. Students develop both theoretical mastery and practical expertise by combining foundational concepts with the construction of influential AI systems.

The curriculum covers foundational materials in search, machine learning, reinforcement learning, and probability. Students then build on these to construct detailed implementations of landmark AI systems such as

AlexNet, GPT-2, and AlphaZero. This rigorous approach develops the analytical skills needed to build the future AI. Finally, as an essential component, this course will address the ethics and responsible development of AI/ML technology and products.

The course emphasizes both technical excellence and ethical considerations in AI development. It serves as the foundation for 07-380 Artificial Intelligence and Machine Learning II, which explores advanced topics, research methods, and specialized applications.

Key topics: AI/ML ethics and responsible development, heuristic-based search, adversarial search, randomized search, constraint satisfaction problems, decision trees, linear/logistic regression, gradient descent-based optimization, neural networks, backpropagation, regularization, cross-validation, pre-training/fine tuning, attention/transformer networks, feature engineering, supervised/unsupervised/self-supervised learning, feature learning, dimensionality reduction, reinforcement learning, approximate/deep RL, maximum (conditional) likelihood estimation for ML

Rationale

07-280, AI & ML I, and the subsequent 07-380, AI & ML II, are a two course sequence designed to replace older AI and ML courses, 15-281 and 10-315, respectively.

Like the previous pair of courses, the 07-280 and 07-380 sequence covers the breadth and depth required by the AI majors. With this redesign, the first of the two courses, 07-280 covers the core AI and ML concepts and provides a great single-course option for any student who wants a good technical introduction to the field of AI.

This restructure will provide the following benefits:

- Flexibility to grow two AI courses
    - Building on first course in the second course
    - Adapting topics more easily across the two course
- Better single AI course for non-AI majors
    - First course as accessible as 15-281 was
    - First course includes core ML topics in addition to AI breadth

Related courses

15-281 and 10-315: Being phased out as part of the 07-280 + 07-380 redesign.

10-301/601: 07-280 overlaps significantly with this flagship intro machine learning course. Specifically, lectures 5-21, 23-24 in the lecture schedule listed below. See course differences section later in this document.

10-701: Similar to 10-301/601, 10-701 will overlap significantly with 07-280 as it is very close in scope to 10-301/601.

Learning Objectives

After completing the course, students should be able to:

- Implement and analyze existing algorithms, including well-studied methods for heuristic-based search, constraint satisfaction problems, classification, regression, and feature learning
- Integrate multiple facets of practical AI and ML in a single system: problem formulation, data preprocessing, learning, regularization, and algorithm/model selection
- Describe the formal properties of algorithms and models in AI and ML, and explain the practical implications of those results
- Describe the trade-offs between computationally cheap, less efficient/accurate methods and computationally expensive, but more accurate methods, especially as they apply to heuristic implementation and sampling techniques
- Compare and contrast different paradigms for learning (supervised, unsupervised, etc.)
- Design experiments to evaluate and compare different machine learning techniques on real-world problems
- Employ probability, statistics, calculus, linear algebra, and optimization in order to develop new predictive models or learning methods
- Given a description of an AI/ML technique, analyze it to identify (1) the expressive power of the formalism; (2) the size and complexity of the search space; (3) the computational properties of the algorithm; (4) any guarantees (or lack thereof) regarding termination, convergence, correctness, accuracy, or generalization power
- Select and apply an appropriate supervised learning algorithm for classification and regression problems (e.g., decision trees, linear regression, logistic regression, neural networks)

Lecture Schedule (subject to change)

Two 80-minute per week plus one weekly 50-minute recitation section.

Example lecture schedule (assuming in-class midterm exams):

Lecture	***Module***	**Lecture Topic(s)**
1	*Intro*	Introduction
2	*Search Fundamentals*	Heuristic Search
3	*Search Fundamentals*	Adversarial Search
4	*Search Fundamentals*	Constraint Satisfaction Problems
5	*ML Fundamentals*	ML Problem Formulation
6	*ML Fundamentals*	Decision Trees
7	*ML Fundamentals*	Linear Regression
8	*ML Fundamentals*	Optimization
9	*ML Fundamentals*	Logistic Regression
10	*ML Fundamentals*	Feature Engineering and Regularization
11	*ML Fundamentals*	Neural networks
12	*ML Fundamentals*	Backpropagation

13	*Building AlexNet*	CNNs and GPUs
14	*Building AlexNet*	Deep Learning Frameworks, Autograd Transfer learning/Fine-tuning
15	*Building AlexNet*	Responsible AI/ML
16	*Building GPT2*	Feature Learning and Word Embedding
17	*Building GPT2*	Tokenization, Position Encoding, Attention
18	*Building GPT2*	Transformer LLMs
19	*RL Fundamentals*	MDPs
20	*RL Fundamentals*	Reinforcement Learning
21	*Building AlphaZero*	Deep RL
22	*Building AlphaZero*	Monte Carlo Tree Search
23	*Probability Fundamentals*	MLE and ML
24	*Probability Fundamentals*	Probabilistic ML Applications
25	*AI / ML Ethics*	AI/ML Ethics

Assessment

Formative assessment:

Students will build their understanding through a series of approximately five written problem sets and approximately five programming assignments. Various rapid feedback questions, such as autograded exercises and in-class polls, will be included as needed.

Summative assessment:

Multiple proctored quizzes and/or exams will be administered to assess students' progress at various intervals throughout the course.