95-891: Introduction to Artificial Intelligence

Fall 2024 schedule (6 units):

- Lectures: TTh 5:00 PM - 6:20 PM Eastern, Hamburg Hall 1002
- Recitations on every other Friday as indicated below 9:30 AM – 10:50 AM Eastern, Hamburg Hall 1002

Instructor: David Steier (steier@andrew.cmu.edu)
Teaching Assistant:

- Prince Wang (princewang@cmu.edu)
- Yubo Li (yubol@andrew.cmu.edu)

Driven by the combination of increased access to data, computational power, and improved sensors and algorithms, artificial intelligence (AI) technologies are entering the mainstream of technological innovation. These technologies include search, computer vision, and natural language processing

The course begins by describing artificial intelligence and provide some examples of AI applications to place developments in the field in historical perspective. We will then do a deep dive into some of the applications underlying search, several varieties of machine learning, computer vision, and natural language processing. The course concludes with a discussion of ethical implications of AI, and a glimpse into potential futures of AI. Exercises will include hands-on application of basic AI techniques as well as selection of appropriate technologies for a given problem and anticipation of design implications. In a final project, groups of students will have the option either of applying natural language processing techniques, or of critically evaluating a recent development in AI (of their choice) by applying published techniques to new data.

Course Learning Outcomes

The main learning objectives of the course are to:

1. Identify problems where artificial intelligence techniques are applicable
2. Apply selected basic AI techniques; judge applicability of more advanced techniques.
3. Participate in the design of systems that act intelligently and learn from experience.

Course Prerequisites

This course is primarily aimed at students with technical backgrounds who wish to design and develop products and services using AI. A background in basic statistics is required for the course. Students need at least a basic knowledge of Python to complete the assignments for this course. Students who have not taken 90-812 or 95-888 or have equivalent background will be required to complete supplementary work to learn Python at the beginning of the course.

Instructor: David Steier (PhD, CMU SCS '89)

David Steier joined the CMU faculty in 2018 as Distinguished Service Professor in the Heinz College School of Information Systems and Management. Prior to joining CMU, David was Managing Director in Deloitte Consulting's Data Science practice. At Deloitte, David helped clients use advanced data analytics and visualization in a variety of industries. Prior to Deloitte, David was Director in the Center for Advanced Research at PwC, Senior Director of Technology and Business Development at Kanisa, and Managing Director at Scient. He was also lead instructor for UC Berkeley's data science capstone course. David holds a Ph.D. in computer science from Carnegie Mellon and a bachelor's degree in computer science from Purdue University.

Course Resources and Policies

Canvas and Piazza

An online site with this syllabus, readings, and other resources has been created in Canvas at https://canvas.cmu.edu/courses/42832 . The system is highly catered to getting you help fast and efficiently from classmates, the TAs, and myself. Rather than emailing questions to the teaching staff, we encourage you to post your questions on Piazza. Find our class signup link at https://piazza.com/cmu/fall2024/95891a1/info .

Text

The primary reference for the course is Russell, S. & Norvig, P. *Artificial Intelligence: A Modern Approach*, Pearson, 2020. This is the fourth edition of the leading textbook in AI, generally accepted as the most comprehensive reference on the subject. It is a substantial update to the third edition, so investing in the latest edition is worthwhile, especially if you plan to do further work in AI. One copy is also available on virtual reserve through the CMU library. This text will be supplemented by on-line material as listed in the course outline below.

Homework, final projects and grading

There will be 2 homework assignments, meant to be done individually, each due at midnight (Eastern Time) per the following schedule:

	Assignment	**Due**
1)	Search	Sep 12

2) Computer vision Sep 26

Each assignment will count for 15 percent of the grade, for 30 percent of the grade. Late assignments (without a written excuse for medical/family/etc. emergencies) will be penalized at the rate of 10% of the assignment's grade per day late. Three closed-book quizzes on Sep 6, Sep 19, and Oct 3 will count for another 30 percent of the grade. There is no final exam for this class. A final project presentation and report will count for 20 percent of the grade. The final projects are meant to be done in groups, which should split the work equally and report who did what in the final report. The remaining 20 percent will be based on class attendance and participation in discussions in the lecture sessions (Tuesdays and Thursdays). In-person attendance is required for attendance credit, unless there are medical circumstances requiring remote attendance. One absence is permitted with further absences causing a proportional deduction in the class participation grade. When you must miss class, please notify me (at least 24 hours in advance except for illness/emergency), so that we can discuss alternative arrangements for catching up on class and associated work. If you encounter extenuating circumstances and must miss more than one class, please come and discuss the issue with me; I would like to find a way to support you. Attendance at the Friday recitation sessions is strongly encouraged they will present helpful information for the assignments, but is not mandatory.

Grading will be on a straight scale as follows (with no rounding up):

A+	98.0-100%	B+	88.0-89.9%	C+	78.0-79.9%
A	92.0-97.9%	B	82.0-87.9%	C	72.0-77.9%
A-	90.0-91.9%	B-	80.0-81.9%	C-	70.0-71.9%

Everyone taking the class should register for a letter grade. Auditing the class, or taking the class Pass/Fail, is intended for extremely rare circumstances and only with consent of the instructor.

Course Outline

This is a mini-course, planned in general around two lecture sessions per week of 1 hour 20 minutes apiece. Lecture sessions will be held in Hamburg Hall 1002. There will also be recitation sessions on the first five Fridays at 9:30 AM – 10:50 AM in Hamburg Hall 1002 as indicated in the schedule below.

- Week 1: Introduction to AI and Search (Lectures Aug 27 & 29, Recitation Aug 30)
  - **Topics**
    - Introductions
    - Course structure and policies
    - What is AI (using case studies)
    - Importance of search for AI
    - Uninformed and informed search
    - Recitation topic: Implementing search
  - **Readings**

    - Russell & Norvig, Chapter 1, “Introduction” and Chapter 3: Solving Problems by Searching,” in *Artificial Intelligence: A Modern Approach*, 2020
- **Assignment out**: HW #1

- Week 2: Adversarial Search and Machine Learning (Lectures Sep 3 & 5, Recitation Sep 6)
    - **Topics**
        - Adversarial search
        - Supervised vs. unsupervised machine learning
        - Regression, classification, clustering
        - Dimensionality reduction
        - Model evaluation
        - Recitation topic: Implementing machine learning
    - **Readings**
        - “Chapter 5: Adversarial Search” in Russell & Norvig, *Artificial Intelligence: A Modern Approach*, 2020
        - Chapter 19: Learning from Examples” in Russell & Norvig, *Artificial Intelligence: A Modern Approach*, 2020
        - C. Aggarwal, Chapter 9 ”Unsupervised Learning” in *Artificial Intelligence: A Textbook*, 2021, Springer, file available on Canvas
    - **Quiz #1** (Sep 6)

- Week 3: Reinforcement and Deep Learning (Lectures Sep 10 & 12, Recitation Sep 13)
    - **Topics**
        - Reinforcement learning
        - Neural networks and back-propagation
        - Convolutional neural networks
        - Recurrent neural networks and LSTMs
        - Transfer learning
        - Recitation: Implementing deep learning in PyTorch
    - **Readings**
        - Russell & Norvig, Chapter 21 & 22 “Deep Learning” and “Reinforcement Learning” in *Artificial Intelligence: A Modern Approach*, 2020
        - (optional) 3Blue1Brown, “But what is a neural network,” Chapter 1 Deep learning,” 2017 (20 min video) https://www.youtube.com/watch?v=aircAruvnKk&vl=en
    - **Assignment due**: HW# 1 (Sep 12)

- Week 4: Computer Vision (Lectures Sep 17 & 19, Recitation Sep 20)
    - **Topics**
        - Introduction to computer vision
        - Image segmentation
        - Edge and motion detection
        - Object classification
        - Pre-trained vision models
        - Recitation topic: Implementing CNNs in PyTorch
    - **Readings**
        - Russell and Norvig, Chapter 25, "Computer Vision" in *Artificial Intelligence: A Modern Approach*, 2020
        - PyTorch Dataloader video: https://www.youtube.com/watch?v=zN49HdDxHi8
        - (optional) TensorFlow, "Image Recognition", July 30, 2018, https://www.tensorflow.org/tutorials/images/image_recognition
    - **Assignment out**: HW #2
    - **Quiz #2** (Sep 19)

- Week 5: Natural Language Processing (Lectures Sep 24 & 26, Recitation Sep 27)
    - **Topics**
        - Intro to natural language understanding
        - Transformers
        - Large language models
        - Retrieval Augmented Generation
        - Vector databases
        - Recitation topic: Implementing natural language processing
    - **Readings**
        - Russell & Norvig, "Chapter 23: Natural Language Processing" and "Chapter 24: Deep Learning for Natural Language Processing," in *Artificial Intelligence: A Modern Approach*, 2020
        - Brown, T.B. et al, "Language Models are Few Shot Learners" (GPT-3), May 28, 2020, https://arxiv.org/pdf/2005.14165.pdf
        - OpenAI, "ChatGPT: Optimizing Language Models for Dialog" November 30, 2022, https://openai.com/blog/chatgpt/
        - L. Monigatti, "A Gentle introduction to Vector Databases, August 1, 2023, https://weaviate.io/blog/what-is-a-vector-database
        - Y. Gao, et. al. "Retrieval-Augmented Generation for Large Language Models," January 4, 2024, https://arxiv.org/abs/2312.10997
        - (optional) Devlin, J. et al. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding", 24 May 2019, https://arxiv.org/pdf/1810.04805.pdf
    - **Assignment due**: HW #2 (Sep 26)
    - **Assignment out**: Final project

- Week 6: AI and Ethics; Futures of AI (Lectures Oct 1 & 3; no recitation)
    - **Topics**
        - Algorithmic bias
        - AI and the future of work
        - Robotics
        - Brain-computer interfaces
        - Quantum computing
        - Artificial general intelligence
    - **Readings**
        - Russell & Norvig, Chapter 27 & 28, " Philosophy, Safety and Ethics of AI"and "The Future of AI" in *Artificial Intelligence: A Modern Approach*, 2020
        - E. Ntoutsi, et. al., "Bias in data-driven artificial intelligence systems—An introductory survey," Wiley Online Library, 03 February 2020, https://onlinelibrary.wiley.com/doi/full/10.1002/widm.1356
        - Russell & Norvig, "Chapter 26.1-26.7: Robotics" in *Artificial Intelligence: A Modern Approach*, 2020
        - J. Reed, et. Al, "A Generalist Agent", May 19, 2022, https://arxiv.org/pdf/2205.06175.pdf
    - **Quiz #3** (Oct 3)
- Week 7: Final Project Presentations (Oct 8 & 10; no recitation)
    - **Topics**
        - Final project presentations
    - **Assignment due**
        - Final project presentations (Oct 8 & 10)
        - Final report (Oct 15)

Academic Integrity

Students are expected to strictly follow Carnegie Mellon University rules of academic integrity in this course. This means in particular that unless otherwise specified, homework are to be the work of the individual student using only permitted material and without any cooperation of other students or third parties. It also means that usage of work by others is only permitted in the form of quotations and any such quotation must be distinctively marked to enable identification of the student's own work and own ideas. All external sources used must be properly cited, including author name(s), publication title, year of publication, and a complete reference needed for retrieval. The same work may not be submitted for credit in multiple courses. Violations will be penalized to the full extent mandated by the CMU policies. There will be no exceptions.

Certain assignments in this course will permit or even encourage the use of generative artificial intelligence (AI) tools, such as ChatGPT. When AI use is permissible, it will be clearly stated in the assignment prompt posted in Canvas. Otherwise, the default is that use of generative AI is disallowed. In assignments where generative AI tools are allowed, their use must be appropriately acknowledged and cited. For instance, if you generated the whole document through ChatGPT and edited it for accuracy, your submitted work would need to include a note such as "I generated this work through Chat GPT and edited the content for accuracy." Paraphrasing or quoting smaller samples of AI generated content must be appropriately acknowledged and cited, following the guidelines established by the APA Style Guide. It is each student's responsibility to assess the validity and applicability of any AI output that is submitted. You may not earn full credit if inaccurate on invalid information is found in your work. Deviations from the guidelines above will be considered violations of CMU's academic integrity policy. Note that expectations for "plagiarism, cheating, and acceptable assistance" on student work may vary across your courses and instructors. Please email me if you have questions regarding what is permissible and not for a particular course or assignment.

Diversity

It is my intent that students from all diverse backgrounds and perspectives be well served by this course, that students' learning needs be addressed both in and out of class, and that the diversity that students bring to this class be viewed as a resource, strength and benefit. It is my intent to present materials and activities that are respectful of diversity: gender, sexuality, disability, age, socioeconomic status, ethnicity, race, and culture. Your suggestions are encouraged and appreciated. Please let me know ways to improve the effectiveness of the course for you personally or for other students or student groups.

Disability Accommodations

If you have a disability and have an accommodations letter from the Disability Resources office, I encourage you to discuss your accommodations and needs with me as early in the semester as possible. I will work with you to ensure that accommodations are provided as appropriate. If you suspect that you may have a disability and would benefit from accommodations but are not yet registered with the Office of Disability Resources, I encourage you to contact them at access@andrew.cmu.edu.

Mental Health

As a student, you may experience a range of challenges that can interfere with learning, such as strained relationships, increased anxiety, substance use, feeling down, difficulty concentrating and/or lack of motivation. These mental health concerns or stressful events may diminish your academic performance and/or reduce your ability to participate in daily activities. CMU services are available, and treatment does work. You can learn more about confidential mental health services available on campus at: http://www.cmu.edu/counseling/. Support is always available (24/7) from Counseling and Psychological Services: 412-268-2922.

