95-891: Introduction to Artificial Intelligence

Spring 2021 (12 units); TTh 5:10 PM - 6:30 PM Eastern, remote only

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

NOTE: The content below is from the Fall 2020 delivery of this course and will be updated for Spring 2021

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, machine learning, natural language processing, robotics and image processing.

The course begins by describing what the latest generation of artificial intelligence techniques can actually do. After an introduction of some basic concepts and techniques, the course illustrates both the potential and current limitations of these techniques with examples from a variety of applications. We spend some time on understanding the strengths and weaknesses of human decision-making and learning, specifically in combination with AI systems and on ethical and policy implications of new AI capabilities. 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 participate in the creation of an AI-based application.

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 including health care, banking, retail, manufacturing, telecommunications, media and the public sector. 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.

In addition to his CMU affiliation, David is also a Lecturer at the University of California Berkeley's School of Information, where he is course lead for the data science capstone class in the Masters in Information and Data Science program. 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

An online site with this syllabus, readings, and other resources has been created in Canvas at https://canvas.cmu.edu/courses/19387. We will use Piazza for questions and answers.

Readings

The primary reference for the course is Russell, S. & Norvig, P. *Artificial Intelligence: A Modern Approach*, Prentice-Hall, 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.
This text will be supplemented by on-line material as listed in the course outline below.

Homework, final projects and grading

There will be 6 homework assignments each due at midnight (Eastern Time) per the following schedule:

	Assignment	**Due**
1)	Search	Sep 17
2)	Classification and clustering	Oct 8
3)	Computer vision	Oct 22
4)	Natural language	Nov 5
5)	AutoML	Nov 24
6)	Fairness in AI	Dec 3

Each assignment will count for 10 percent of the grade, with the lowest grade dropped, for a total of 50%. 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. A final project, presented on the final Thursday of the class (Dec 10) with a final report due the following Tuesday (Dec 15), will count for 20 percent of the grade. Occasional quizzes will count for another 10% of the grade. The remaining 20 percent will be based on class attendance and participation in discussions. Two absences are permitted, with further absences causing a proportional deduction in the class participation grade. There is no final exam for this class.

Grading will be on a straight scale as follows:

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 expect to 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.

All synchronous classes will be recorded via Zoom so that students in this course (and only students in this course) can watch or re-watch past class sessions. Please note that breakout rooms will not be recorded. I will make the recordings available on Canvas as soon as possible after each class session (usually within 3 hours of the class meeting). Recordings will live in our Canvas website https://canvas.cmu.edu/courses/19387. Please note that you are not allowed to share these recordings. This is to protect your FERPA rights and those of your fellow students.

Course Outline

This is a full-semester course, planned in general around two sessions per week of 1 hour 20 minutes apiece. Note that due to the Thanksgiving holiday, there is no class on Nov 26. The course will be taught entirely on Zoom.

- Week 1: Introduction to AI (Sep 1 and 3)
  - **Topics**
    - Introduction to artificial intelligence
    - Course structure and policies
    - History of AI
    - Proposing and evaluating AI applications
    - Case study: Google Duplex
  - **Readings**
    - Russell & Norvig, "Chapter 1: Introduction" in *Artificial Intelligence: A Modern Approach*, 2020
    - Amadeo, R., June 27, 2018, "Talking to Google Duplex: Google's human-like phone AI feels revolutionary" https://arstechnica.com/gadgets/2018/06/google-duplex-is-calling-we-talk-to-the-revolutionary-but-limited-phone-ai/
  - **Assignment out**: HW #1

- Week 2: Search (Sep 8 & 10)
  - **Topics**
    - Problem spaces and search
    - Knowledge and rationality
    - Heuristic search strategies

        - Search and optimization (gradient descent)
        - Adversarial search
    - **Readings**
        - "Chapters 3: Solving Problems by Searching," and "Chapter 5: Adversarial search" in Russell & Norvig, *Artificial Intelligence: A Modern Approach*, 2020

- Week 3: Reasoning with Uncertainty and Time (Sep 15 & 17)
    - **Topics**
        - Uncertainty
        - Bayesian reasoning
        - Hidden Markov Models
    - **Readings** in Russell & Norvig, *Artificial Intelligence: A Modern Approach*, 2020
        - Quantifying uncertainty: Ch. 12
        - Probabilistic reasoning: Ch. 13.1-13.3
        - Probabilistic reasoning over time: 14.1-14.3
    - **Assignment due**: HW# 1 (Sep 17)

- Week 4: Machine learning: Supervised methods (Sep 22 & 24)
    - **Topics**
        - What is machine learning?
        - Supervised vs. unsupervised learning
        - Regression -- linear, logistic, ridge
        - Classification – decision trees, SVM, random forests
        - Model performance evaluation – MSE, lift, AUC, Type 1 vs 2 errors
    - **Readings**
        - "Chapter 19: Learning from Examples" in Russell & Norvig, *Artificial Intelligence: A Modern Approach*, 2020
        - (optional) Chapter 5.1-5.7, "ML Basics" in Goodfellow, I., Bengio, Y. and Courville, A. *Deep Learning*, MIT Press, 2016. https://www.deeplearningbook.org/ .
    - **Assignment out**: HW #2

- Week 5: Machine Learning: Unsupervised, Reinforcement and Semi-supervised Learning (Sep 29 & Oct 1)
    - **Topics**
        - Dimensionality reduction: PCA
        - Clustering – k-means, hierarchical clustering
        - Reinforcement learning
        - Semi-supervised learning
    - **Readings**

        - Chapter 5.8: "Unsupervised Machine Learning," in Goodfellow, I., Bengio, Y. and Courville A., Deep Learning, 2016. https://www.deeplearningbook.org/ .
        - Huneycutt, J., "An introduction to clustering algorithms in Python", May 29, 2018, https://towardsdatascience.com/an-introduction-to-clustering-algorithms-in-python-12343857409 7
        - Russell & Norvig, Chapter 22 "Reinforcement Learning" in *Artificial Intelligence: A Modern Approach*, 2020

- Week 6: Deep Learning (Oct 6 & 8)
    - **Topics**
        - Neural networks and back-propagation
        - Convolutional neural networks
        - Recurrent neural networks and LSTMs
        - Transfer learning
    - **Readings:**
        - Chapter 1 " Using neural networks to recognize handwritten digits," in Nielsen, M. A., *Neural Networks and Deep Learning*, Determination Press, 2015, available at http://neuralnetworksanddeeplearning.com/chap1.html
        - Russell & Norvig, Chapter 21 "Deep Learning" in *Artificial Intelligence: A Modern Approach*, 2020
        - Visualizing neural networks using the TensorFlow Playground: https://playground.tensorflow.org/
        - 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 #2 (Oct 8)
    - **Assignment out**: HW #3

- Week 7: Computer Vision (October 13 & 15)
    - **Topics**
        - Introduction to computer vision
        - Image segmentation
        - Object and motion detection
        - Object classification
    - **Readings**
        - "Chapter 25: Computer Vision" in Russell and Norvig, *Artificial Intelligence: A Modern Approach*, 2020
        - Visualizing CNNs, http://scs.ryerson.ca/~aharley/vis/conv/flat.html
        - PyTorch Dataloader video: https://www.youtube.com/watch?v=zN49HdDxHi8
        - TensorFlow, "Image Recognition", July 30, 2018, https://www.tensorflow.org/tutorials/images/image_recognition
    - **Assignment due**: Final project groups formed and initial proposal (Oct 15)

- Week 8: Natural Language Understanding (Oct 20 & 22)
    - **Topics**
        - Intro to natural language understanding
        - Machine translation
        - Sentiment analysis
        - Application of deep learning to NLP
    - **Readings**
        - Lewis-Krause, G. “The Great AI Awakening” , The New York Times,, December 14, 2016, https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html
        - Russell & Norvig, “Chapter 23: Natural Language Processing” and “Chapter 24: Deep Learning for Natural Language Processing,” in *Artificial Intelligence: A Modern Approach*, 2020
        - (Optional) Brown, T.B. et al, “Language Models are Few Shot Learners” (GPT-3), May 28, 2020, https://arxiv.org/pdf/2005.14165.pdf
    - **Assignment due**: HW #3 (Oct 22)
    - **Assignment out**: HW #4

- Week 9: Speech and Natural Language Interaction (Oct 27 & 29)
    - **Topics**
        - Speech recognition
        - Speech synthesis
        - Natural language generation
        - Chatbots
        - Case study: Google Duplex (revisited)
    - **Readings**
        - Leviathan Y, and Matias, Y, “Google Duplex: An AI System for Accomplishing Real-World Tasks Over the Phone,” Google AI Blog, May 8, 2018 https://ai.googleblog.com/2018/05/duplex-ai-system-for-natural-conversation.html
        - Daly, L. “Chatbot Fundamentals: An interactive guide to writing bots in Python”, 2016 https://apps.worldwritable.com/tutorials/chatbot/

- Week 10: Robotics (Nov 3 & 5)
    - **Topics**
        - Introduction to robotics
        - Human-robot interaction
        - Navigation and path planning
        - Learning and robotics: Reinforcement learning
        - Autonomous vehicles technologies and impacts
    - **Readings**

    - Russell & Norvig, "Chapter 26: Robotics" in *Artificial Intelligence: A Modern Approach*, 2020
    - Boston Dynamics videos, 2018 , https://www.youtube.com/user/BostonDynamics
    - Priday, R. "What's Really Going on in those Boston Dynamics Videos," *Wired*, February 18, 2018 https://www.wired.co.uk/article/boston-dynamics-robotics-roboticist-how-to-watch
- **Assignment due**: HW #4 (Nov 5)
- **Assignment out**: HW #5

- Week 11: Explainability and Automation of AI/ML (Nov 10 & 12)
    - **Topics**
        - Explainability and interpretability of AI models
        - AutoML
    - **Readings**
        - J. Zornoza, "Explainable Artificial Intelligence" April 15, 2020, https://towardsdatascience.com/explainable-artificial-intelligence-14944563cc79
        - N. Erickson, "AutoGluon-Tabular: Robust and Accurate AutoML for Structured Data," 13 Mar 2020, https://arxiv.org/pdf/2003.06505.pdf

- Week 12: Ethical and Legal Considerations in AI (Nov 17 & 19)
    - **Topics**
        - Algorithmic bias
        - AI and the future of work
        - Privacy
        - Appropriate uses of AI
    - **Readings**
        - Russell & Norvig, "Chapter 27: Philosophy, Safety and Ethics of AI" in *Artificial Intelligence: A Modern Approach*, 2020
        - Erik Brynjolfsson, Tom Mitchell, and Daniel Rock, "What Can Machines Learn and What Does It Mean for Occupations and the Economy?", AEA Papers and Proceedings 2018, 108: 43–47, http://ide.mit.edu/sites/default/files/publications/pandp.20181019.pdf
        - R. "Bias detectives: the researchers striving to make algorithms fair," *Nature*, June 2018, https://www.nature.com/magazine-assets/d41586-018-05469-3/d41586-018-05469-3.pdf

- Week 13: Emotion and the Arts in AI (Nov 24, no class Nov 26)
    - **Topics**
        - Using AI to recognize emotion
        - AI in writing, music, and art
    - **Readings**

        - J. Sukis The Relationship Between Art and AI," May 20, 2018, https://medium.com/design-ibm/the-role-of-art-in-ai-31033ad7c54e#:~:text=We%20began%20using%20AI%20to,and%20applies%20them%20to%20another.
    - **Assignment out**: HW #6
    - **Assignment due**: HW #5 (Nov 24)

- Week 14: Infrastructure for AI (Dec 1 & 3)
    - **Topics**
        - Parallel and distributed computing for scalability:
        - Resolving technical tradeoffs
        - Data pipelines
    - **Readings**
        - Robinson, J. "How Facebook Scales Machine Learning," Feb 3, 2019, https://medium.com/@jamal.robinson/how-facebook-scales-artificial-intelligence-machine-learning-693706ae296f
        - V. Paruchuri, "Building An Analytics Data Pipeline In Python," March 15, 2017, https://www.dataquest.io/blog/data-pipelines-tutorial/
        - Zheng, H. Wang, Y, and Molino, P. "COTA: Improving Uber Customer Care with NLP & Machine Learning," January 2018, https://eng.uber.com/cota/
    - **Assignment due**: HW #6 (Dec 3)

- Week 15: The Future of AI and Final Project Presentations (Dec 8 and 10)
    - **Topics**
        - Emerging developments like brain-computer interfaces
        - Final project presentations and wrap-up
    - **Readings**
        - Russell & Norvig, "Chapter 28: The Future of AI", in *Artificial Intelligence: A Modern Approach*, 2020
    - **Assignment due**
        - Final project presentations (Dec 10)
        - Final project reports due (Dec 15)

Academic Integrity

Academic Integrity is expected at all time. Carnegie Mellon has an established well-defined policy on this subject which can be found at:
http://www.cmu.edu/policies/documents/Academic%20Integrity.htm
It is the responsibility of the student to verse themselves with these policies. All necessary and appropriate sanctions will be issued to all parties involved with plagiarizing any and all course

work. Assignments are expected to be done individually, with the exception of the final projects. You must acknowledge in your submission any help received on your assignments. That is, you must include a comment in your homework submission that clearly states the name of the student, book, or online reference from which you received assistance. External sources should be appropriately credited in assignments, presentations and reports. Plagiarism and any other form of academic dishonesty that is in violation with these policies will not be tolerated. In particular, the same work may not be submitted for credit in multiple courses.

Maintaining a Healthy Balance (as recommended and supported by the university)

Do your best to maintain a healthy lifestyle this semester by eating well, exercising, avoiding drugs and alcohol, getting enough sleep and taking some time to relax. This will help you achieve your goals and cope with stress. All of us benefit from support during times of struggle.

You are not alone. There are many helpful resources available on campus and an important part of the college experience is learning how to ask for help. Asking for support sooner rather than later is often helpful. If you or anyone you know experiences any academic stress, difficult life events, or feelings like anxiety or depression, we strongly encourage you to seek support.

For local help and referrals, please contact the Office of the Dean of Student Affairs in 301 Warner Hall (412-268-2075). Counseling and Psychological Services (CaPS) at the Pittsburgh campus can also help you get connected to support. You can call them at 412-268-2922 and/or visit their website at http://www.cmu.edu/counseling/ to learn more.

If you or someone you know is feeling suicidal or in danger of self-harm, call someone immediately, day or night:

CaPS: 412-268-2922
Suicide Prevention Hotline: 800-273-8255 (TALK)

