95-891: Introduction to Artificial Intelligence

Fall 2021 (12 units); TTh 4:40 PM - 6:00 PM Eastern, Hamburg Hall 2009

Instructor: David Steier (steier@andrew.cmu.edu)
Teaching Assistant: Yue Zhao (zhaoy@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, 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 and Piazza

An online site with this syllabus, readings, and other resources has been created in Canvas at https://canvas.cmu.edu/courses/25164 . This term we will be using Piazza for class discussion. The system is highly catered to getting you help fast and efficiently from classmates, the TA, 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 piazza.com/cmu/fall2021/95891a

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. 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 16
2)	Classification and clustering	Oct 7
3)	Computer vision	Oct 30
4)	Natural language	Nov 4
5)	AutoML	Nov 18
6)	Fairness in AI	Dec 2

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 presentation and report will count for 20 percent of the grade. Three 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.

Course Outline

This is a full-semester course, planned in general around two sessions per week of 1 hour 20 minutes apiece. Note that there is no class on October 14 due to mid-semester break and no class on November 25 due to the Thanksgiving holiday.

- Week 1: Introduction to AI (Aug 31 & Sep 2)
  - **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., "Talking to Google Duplex: Google's human-like phone AI feels revolutionary", June 27, 2018, 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 7 & 9)
  - **Topics**
    - Importance of search for AI
    - Uninformed and informed search
    - Adversarial search
    - Local search (gradient descent)
  - **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 (Sep 14 & 16)
  - **Topics**

        - Uncertainty
        - Bayesian networks
    - **Readings** in Russell & Norvig, *Artificial Intelligence: A Modern Approach*, 2020
        - Quantifying uncertainty: Ch. 12
        - Probabilistic reasoning: Ch. 13.1-13.3
    - **Assignment due**: HW# 1 (Sep 16)

- Week 4: Introduction to machine learning (Sep 21 & 23)
    - **Topics**
        - What is machine learning?
        - Supervised vs. unsupervised learning
        - Regression -- linear, logistic, ridge
        - Classification – decision trees, SVM, random forests
        - Model evaluation
        - Dimensionality reduction: PCA
        - Clustering – k-means, hierarchical clustering
    - **Readings**
        - "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
        - Huneycutt, J., "An introduction to clustering algorithms in Python", May 29, 2018, https://towardsdatascience.com/an-introduction-to-clustering-algorithms-in-python-123438574097
        - (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: Probabilistic and reinforcement learning (Sep 28 & 30)
    - **Topics**
        - Bayesian learning
        - Expectation maximization
        - Reinforcement learning
    - **Readings**
        - Russell & Norvig, Chapters 20 and 22, "Learning Probabilistic Models" and "Reinforcement Learning" in *Artificial Intelligence: A Modern Approach*, 2020
        - (optional) van Engelen, J.E., Hoos, H.H. A survey on semi-supervised learning. Mach Learn 109, 373–440 (2020). https://link.springer.com/article/10.1007/s10994-019-05855-6

- Week 6: Deep Learning (Oct 5 & 7)
    - **Topics**
        - Neural networks and back-propagation
        - Convolutional neural networks
        - Recurrent neural networks and LSTMs
        - Transfer learning
    - **Readings:**
        - Russell & Norvig, Chapter 21, "Deep Learning" in *Artificial Intelligence: A Modern Approach*, 2020
        - Nielsen, M. A., Chapter 1 " Using neural networks to recognize handwritten digits," in *Neural Networks and Deep Learning*, Determination Press, 2015, available at http://neuralnetworksanddeeplearning.com/chap1.html
        - 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 7)
    - **Assignment out**: HW #3

- Week 7: Computer Vision (Oct 12, no class Oct 14)
    - **Topics**
        - Introduction to computer vision
        - Image segmentation
        - Object and motion detection
        - Object classification
    - **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
        - TensorFlow, "Image Recognition", July 30, 2018, https://www.tensorflow.org/tutorials/images/image_recognition
    - **Assignment due**: Final project groups formed and initial proposal (Oct 12)

- Week 8: Natural Language Understanding (Oct 19 & 21)
    - **Topics**
        - Intro to natural language understanding
        - Machine translation
        - Sentiment analysis
        - Language models
    - **Readings**

    - 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

- Week 9: Speech and Natural Language Interaction (Oct 26 & 28)
  - **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/
  - **Assignment due**: HW #3 (Oct 30)
  - **Assignment out**: HW #4

- Week 10: AutoML; Introduction to Robotics (Nov 2 & 4)
  - **Topics**
    - AutoML
    - Introduction to robotics
    - Navigation and path planning
    - Learning and robotics: Reinforcement learning
  - **Readings**
    - Hutter, Korrhoffm, and Vanschoren, "Chapter 1: Automated Machine Learning", 2020, https://library.oapen.org/bitstream/handle/20.500.12657/23012/1007149.pdf
    - N. Erickson, "AutoGluon-Tabular: Robust and Accurate AutoML for Structured Data," 13 Mar 2020, https://arxiv.org/pdf/2003.06505.pdf
    - Russell & Norvig, "Chapter 26.1-26.7: Robotics" in *Artificial Intelligence: A Modern Approach*, 2020
    - Boston Dynamics videos https://www.youtube.com/user/BostonDynamics
    - Adam Savage, "How Boston Dynamics' Spot Robot Works," Apr 8, 2020, (19 minute video), https://www.youtube.com/watch?v=R-PdPtqw78k&t=896s
  - **Assignment due**: HW #4 (Nov 7)

- **Assignment out**: HW #5

- Week 11: Human-robot interaction and autonomous vehicles (Nov 9 & 11)
  - **Topics**
    - Human-robot interaction
    - Approaches to autonomy
    - Autonomous vehicles technologies and impacts
  - **Readings**
    - Russell & Norvig, "Chapter 26.8 Humans and Robots" in *Artificial Intelligence: A Modern Approach*, 2020
    - L. Fridman, "Human-Centered Autonomous Vehicle Systems: Principles of Effective Shared Autonomy", 3 Oct 2018, https://arxiv.org/pdf/1810.01835.pdf
    - Ali Nasseri, et. al, "Autonomous Vehicle Technology Report, 2020", https://www.wevolver.com/article/2020.autonomous.vehicle.technology.report

- Week 12: Explainability; AI & Ethics (Nov 16 & 18)
  - **Topics**
    - Explainability and interpretability of AI models
    - Algorithmic bias
    - AI and the future of work
    - Privacy
    - Appropriate uses of AI
  - **Readings**
    - J. Zornoza, "Explainable Artificial Intelligence" April 15, 2020, https://towardsdatascience.com/explainable-artificial-intelligence-14944563cc79
    - 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
  - **Assignment due**: HW #5 (Nov 18)
  - **Assignment out**: HW #6

- Week 13: The Future of AI (Nov 23, no class on Nov 25)
  - **Topics**
    - AI in writing, music, and art
    - Emerging developments like brain-computer interfaces

- **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.
    - Russell & Norvig, "Chapter 28: The Future of AI", in *Artificial Intelligence: A Modern Approach*, 2020

- Week 14: Infrastructure for AI; Final project presentations (Nov 30 & Dec 2)
    - **Topics**
        - Parallel and distributed computing for scalability
        - MLOps
        - MLSys
    - **Readings**
        - https://a16z.com/2020/10/15/the-emerging-architectures-for-modern-data-infrastructure/
        - Google Cloud, MLOps: Continuous delivery and automation pipelines in machine learning, https://cloud.google.com/architecture/mlops-continuous-delivery-and-automation-pipelines-in-machine-learning
        - Robinson, J. "How Facebook Scales Machine Learning," Feb 3, 2019, https://medium.com/@jamal.robinson/how-facebook-scales-artificial-intelligence-machine-learning-693706ae296f
        - Y. Zhao, et. al., "Suod: Accelerating Large-Scale Unsupervised Heterogeneous Outlier Detection," in *MLSys '21*, https://www.andrew.cmu.edu/user/yuezhao2/papers/21-mlsys-suod.pdf
    - **Assignment due**:
        - HW #6 (Dec 2)
        - Final project presentations, round 1 (Dec 2)

- Week 15: Final Project Presentations, continued (Dec 7)
    - **Topics**
        - Emerging developments like brain-computer interfaces
        - Final project presentations and wrap-up
    - **Readings**
    - **Assignment due**
        - Final project presentations, round 2 (Dec 7)
        - Final project reports due (Dec 9)

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.

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.

