95-891: Introduction to Artificial Intelligence - Full Syllabus

Spring 2019 (12 units)

Class times: T/Th 4:30 pm - 5:50 pm (Eastern time), Hamburg Hall 1204 and on-line
Instructor: David Steier (steier@andrew.cmu.edu
Teaching assistant: Abhinav Maurya (amaurya@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, 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. 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. Programming experience is also required as the course will includes hands-on programming exercises to illustrate the concepts.

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 co-instructs 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.

The Teaching Assistant for this class will be Abhinav Maurya, who is a PhD student in the Heinz school, specializing in machine learning and public policy.

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/8663.

Readings

There are two primary reference for the course:

- Russell, S. & Norvig, P *Artificial Intelligence: A Modern Approach*, Prentice-Hall, 2010. This is the third edition of the leading textbook in AI, generally accepted as the most comprehensive reference to AI. However, some of the content is by now somewhat dated, and we will mainly use selected chapters as references on foundational content.
- Goodfellow, I., Bengio, Y. and Courville, A., *Deep Learning*, MIT Press, 2016. Includes more modern material on machine learning and deep learning. May be purchased but is also freely available from https://www.deeplearningbook.org/ .

These texts 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)	Problem spaces and search	Jan 31
2)	Classification and clustering	Feb 19
3)	Image classification	Mar 7
4)	Natural language	Mar 28
5)	Robotics	Apr 9
6)	Fairness in AI	Apr 25

Each assignment will count for 10 percent of the grade. A final project, presented on the final Thursday of the class (May 2) with a final report due the following Tuesday (May 7), will count for 20 percent of the grade. The remaining 20 percent will be class participation: class attendance and participation in discussions. There is no final exam for this class.

Grading will be on a straight scale as follows:

A+	98-100%	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 will receive a letter grade; auditing the class, or taking the class Pass/Fail, is not permitted.

Course Outline

This is a full-semester course, planned in general around two sessions per week of 1 hour 20 minutes apiece; however, note that due to CMU holidays, there will be no class March 12th , 14th and April 11th. Per the schedule below, 15 sessions of the course will be taught in Hamburg Hall 1204, and the remaining 13 sessions will be taught online using Zoom.

- Week 1: Introduction to AI (Jan 15 & 17, HbH 1204)
  - **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*, 2010
    - Chui, et. al, "Notes From the AI Frontier," McKinsey Global Institute, April 2018, available at https://www.mckinsey.com/~/media/mckinsey/featured%20insights/artificial%20intelligence/notes%20from%20the%20ai%20frontier%20applications%20and%20value%20of%20deep%20learning/mgi_notes-from-ai-frontier_discussion-paper.ashx
    - 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 and Planning (Jan 22 & 24, HbH 1204)
  - **Topics**
    - Problem spaces and search
    - Knowledge and rationality
    - Heuristic search strategies
    - Search and optimization (gradient descent)
    - Adversarial search
    - Planning and scheduling
    - Case studies: Playing chess, Manufacturing scheduling

- **Readings**
    - "Chapters 3: Solving Problems by Searching," "Chapter 5: Adversarial search", "Chapter 10.2-10.5: Planning", "Chapter 11: Planning and Acting in the Real World" in Russell & Norvig, *Artificial Intelligence: A Modern Approach*, 2010

- Week 3: Knowledge Representation and Reasoning (Jan 29 & 31, Online)
    - **Topics**
        - Logic and inference
        - Ontologies
        - Bayesian reasoning
        - Temporal reasoning
        - Case study: Medical diagnosis
    - **Readings** in Russell & Norvig, *Artificial Intelligence: A Modern Approach*, 2010
        - Propositional logic: Sections 7.3-7.6
        - First-order logic: Sections 8.1.2 - 8.2, 8.4, 9.1-9.5
        - Knowledge representation: Sections 12.1-12.5, 12.7
        - Quantifying uncertainty: Ch. 13
        - Probabilistic reasoning: Ch. 14.1-14.2
        - (Optional) Probabilistic reasoning over time: Ch. 15
    - **Assignment due**: HW# 1 (Jan 31)

- Week 4: Machine learning: Supervised methods (Feb 5, HbH 1204 & Feb 7, online)
    - **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
        - Case study: Bank failure prediction (also part of HW#2)
    - **Readings**
        - Chapter 18 in Russell & Norvig, *Artificial Intelligence: A Modern Approach*, 2010
        - Chapter 5.1-5.7, "ML Basics" in Goodfellow, I., Bengio, Y. and Courville A., *Deep Learning*, 2016.
    - **Assignment out**: HW #2

- Week 5: Machine Learning: Unsupervised Methods (Feb 12 & 14, HbH 1204)
    - **Topics**
        - Dimensionality reduction: PCA
        - Clustering – k-means, hierarchical clustering
        - Semi-supervised methods
        - Reinforcement learning
        - Choosing among machine learning techniques
        - Case study: Public health outcome clustering
    - **Readings**
        - Chapter 5.8: "Unsupervised Machine Learning," in Goodfellow, I., Bengio, Y. and Courville A., Deep Learning, 2016.
        - 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 21 "Reinforcement Learning" in *Artificial Intelligence: A Modern Approach*, 2010
        - Booz Allen Hamilton "Guide to Analytic Selection," p. 65-83 in *The Field Guide to Data Science*, 2nd Edition 2015, available online at https://www.boozallen.com/content/dam/boozallen_site/sig/pdf/publications/2015-field-guide-to-data-science.pdf

- Week 6: Deep Learning (Feb 19 & 21 HbH 1204)
    - **Topics**
        - Neural networks and back-propagation
        - Convolutional neural networks
        - Recurrent neural networks and LSTMs
    - **Readings:**
        - Chapter 1 " Using neural netowrks to recognize handwritten digits," in Nielsen, M. A., Neural Networks and Deep Learning, Determination Press, 2015, available at http://neuralnetworksanddeeplearning.com/chap1.html
        - Chapter 9 "Convolutional Networks, and Chapter 10, "Sequence Modeling: Recurrent and Recursive Nets" in Goodfellow, I., Bengio, Y. and Courville A., *Deep Learning*, 2016.
        - Visualizing neural networks using the TensorFlow Playground: https://playground.tensorflow.org/
        - (Optional) Chapter 7, "Regularization in Deep Learning," and Chapter 8, "Optimization for Training Deep Models" in Goodfellow, I., Bengio, Y. and Courville A., *Deep Learning*, 2016.
        - (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 #2 (Feb 19)
- **Assignment out**: HW #3

- Week 7: Image Processing (Feb 26 & 28, HbH 1204)
  - **Topics**
    - Introduction to computer vision
    - Image segmentation
    - Object and motion detection
    - Object classification
    - Use of pre-trained models (VGG16, Inception)
  - **Readings**
    - Chapter 24 "Perception" in Russell and Norvig, *Artificial Intelligence: A Modern Approach*, 2010
    - Karpathy, A. "Convolutional Neural Networks for Visual Recognition," 2018, available at http://cs231n.github.io/convolutional-networks/
    - 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

- Week 8: Natural Language Understanding (Mar 5 & 7, HbH 1204)
  - **Topics**
    - Intro to natural language understanding
    - Case study: 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 22: Natural Language Processing" in *Artificial Intelligence: A Modern Approach*, 2010
    - Collobert et al. "Natural Language Processing (Almost) from Scratch," Journal of Machine Learning Research, 2011 available at https://arxiv.org/pdf/1103.0398.pdf
    - (Optional) G Golderg, Y. *Neural Network Methods for Natural Language Processing* Synt*hesis Lectures on Human Language Technologies,* April 2017, freely available monograph at https://doi.org/10.2200/S00762ED1V01Y201703HLT037
    - (Optional) Feldman, R, "Sentiment Analysis Tutorial, IJCAI-13, 2013, http://ijcai13.org/files/tutorial_slides/tf4.pdf
  - **Assignment due**: HW #3 (Mar 7)
  - **Assignment out**: HW #4

- Week 9: Natural Language Interaction (Mar 19 & 21, online)
    - **Topics**
        - Speech recognition
        - Hidden Markov Models
        - Chatbots
        - Natural language generation
        - Speech synthesis
        - Case study: Google Duplex (revisited)
    - **Readings**
        - Russell & Norvig, "Chapter 15.3: Hidden Markov Models" and "Chapter 22: Natural Language for Communication" in *Artificial Intelligence: A Modern Approach*, 2010
        - 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
        - Maruti Labs, "Complete Guide to Chatbots: Development to Promotion," https://www.marutitech.com/complete-guide-chatbots/ 2016
        - Daly, L. "Chatbot Fundamentals: An interactive guide to writing bots in Python", 2016 https://apps.worldwritable.com/tutorials/chatbot/

- Week 10: Robotic Sensing and Manipulation (Mar 26 & 28, Online)
    - **Topics**
        - Introduction to robotics
        - Sensing
        - Manipulation
        - Human-robot interaction
    - **Readings**
        - Russell & Norvig, "Chapter 25: Robotics" in *Artificial Intelligence: A Modern Approach*, 2010
        - Goodrich and Schultz, "Human Robot Interaction: A Survey," 2007 https://faculty.cs.byu.edu/~mike/mikeg/papers/HRISurvey.pdf
    - **Assignment due**: HW #4 (Mar 28)
    - **Assignment out**: HW #5

- Week 11: Mobile Robots (Apr 2 & 4, Online)
    - **Topics**
        - Navigation and path planning
        - Learning and robotics: Reinforcement learning
        - Case study: Autonomous vehicles technologies and impacts
    - **Readings**
        - 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
- RAND Institute, “Autonomous Vehicle Technology: A Guide for Policymakers”, 2016 https://www.rand.org/content/dam/rand/pubs/research_reports/RR400/RR443-2/RAND_RR443-2.pdf

- Week 12: AI in the Enterprise (Apr 9, Online)
  - **Topics**
    - AI in the enterprise
  - **Readings**
    - Beyer, D, “AI and Machine learning in industry,” 2017, download from http://www.oreilly.com/data/free/ai-machine-learning-in-industry.csp?cmp=tw-data-free-article-lgen_tw_free_ebook_as
  - **Assignment due**: HW #5 (Apr 9)
  - **Assignment out**: HW #6

- Week 13: Ethical and Legal Considerations in AI (Apr 16 & 18, Online)
  - **Topics**
    - Privacy
    - Bias
    - AI and the future of work
    - Appropriate uses of AI
    - Case study: AI to predict re-arrests
  - **Readings**
    - Jerome, J, “Why AI may be the next big privacy trend,” https://iapp.org/news/a/why-artificial-intelligence-may-be-the-next-big-privacy-trend/, 2016
    - Burt, A . “How will the GDPR impact machine learning?”, May 16, 2018, “https://www.oreilly.com/ideas/how-will-the-gdpr-impact-machine-learning
    - Vanian, J “Unmasking A.I.'s Bias Problem,” Fortune, June 25, 2018, http://fortune.com/longform/ai-bias-problem/
    - NSTC, “Preparing for the Future of AI,” October 2016
    - Brynjolfsson, E and Mitchell, T. “What can machine learning do? Workforce implications,” *Science* 22 Dec 2017: Vol. 358, Issue 6370, pp. 1530-1534 DOI: 10.1126/science.aap8062
    - Courtland, 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

    - (optional) Abadi, M. et. al "Deep Learning with Differential Privacy, " https://arxiv.org/pdf/1607.00133.pdf

- Week 14: Infrastructure for AI (Apr 23 & 25, Online)
    - **Topics**
        - Parallel and distributed computing for scalability:
        - Resolving technical tradeoffs
        - Case study: Uber and Facebook
    - **Readings**
        - Skymind, "An Introduction to Distributed Training of Neural Networks," November 30, 2017 https://blog.skymind.ai/distributed-deep-learning-part-1-an-introduction-to-distributed-training-of-neural-networks/
        - Levine, S. "Deep Learning for Robots: Learning from Large-Scale Interaction," March 2016, https://ai.googleblog.com/2016/03/deep-learning-for-robots-learning-from.html
        - Zheng, H. Wang, Y, and Molino, P. "COTA: Improving Uber Customer Care with NLP & Machine Learning," January 2018, https://eng.uber.com/cota/
        - Hermann and Del Balso, 2017, "Meet Michelangelo: Uber's Machine Learning Platform," https://eng.uber.com/michelangelo/
        - (Optional) Large Scale Distributed Deep Networks, Jeff Dean et al. (first paper to use tens of thousands of CPU cores for deep learning)
        - (Optional) Deep learning with COTS HPC systems, Adam Coates et al. (first paper to use GPUs for deep learning)
    - **Assignment due**: HW #6 (Apr 25)

- Week 15: The Future of AI & Final Project Presentations (Apr 30 & May 2, HbH 1204)
    - **Topics**
        - The future of AI: Emerging developments
        - Final project presentations and wrap-up
    - **Readings**
        - National Science and Technology Council, "Preparing for the future of AI," October 2016, https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf
    - **Assignment due**
        - Final project presentations (May 2)
        - Final project reports (May 7)

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. 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)

