Department of Psychology – Columbia University

Machine Intelligence

GU4236 - Fall 2020

4 points

Instructor: Trenton Jerde, Ph.D.

Class Meets Online: Wednesdays: 6:10 – 8:00 PM

Email: taj2128@columbia.edu

Office Hours: Online, by appointment

Course bulletin description

This seminar will survey historical and modern developments in machine intelligence from fields such as psychology, neuroscience, and computer science, and from approaches such as cybernetics, artificial intelligence, robotics, connectionism, neural networks, machine learning, and deep learning. The emphasis is on the conceptual understanding of topics. The course does not include, nor require background in, computer programming and statistics. An overall goal is for students to become informed consumers of applications of artificial intelligence.

Prerequisites

Undergraduate students: Undergraduate students should have taken PSYC UN1001 The Science of Psychology or an equivalent introductory psychology course. In addition, some prior coursework or experience in statistics and research methods, as well as cognitive psychology or cognitive neuroscience, would be very helpful. The instructor's permission is required.

Graduate students: Open to Ph.D. students in the psychology department and graduate students in related departments such as computer science, neuroscience, and engineering, with instructor's permission.

Instructor

The instructor is a Senior Editor at *Nature Machine Intelligence*, a scientific journal on artificial intelligence, machine learning, and robotics. He has worked as a consultant on brain disorders, including at Medtronic on deep brain stimulation for Parkinson's disease; and held faculty positions at Columbia University, New York University, and the University of Minnesota. His research interests include machine learning, systems and cognitive neuroscience, and cognitive science. He has a PhD in neuroscience from the University of Minnesota and a BA in psychology with minors in English and philosophy from the University of Iowa.

Detailed description of the course

The goal of the seminar is to achieve a deeper understanding of machine intelligence.

Students will participate in group discussions during class, and give short presentations on current topics of their choice. Students should try to engage the class as much as possible – remember, seminars are meant to be discussions.

Specifically:

- We will ask fundamental questions, such as: What is intelligence? Can machines think? Can machine intelligence approach the human level (artificial general intelligence) or surpass it (superintelligence)?
- We will distinguish data, algorithms, compute, and technology in applications of machine intelligence.
- We will discuss the nature of learning and knowledge acquisition in artificial systems, such as data-driven learning and innate or built-in knowledge.
- We will compare intelligence as manifested in computer programs, simulated agents, bots, drones, robots, etc.
- We will explore how basic psychological and behavioral processes, such as perception, learning, language, thought, emotion, and movement, are embedded in machines.
- We will examine applications of machine intelligence in society, such as facial recognition, and consider the roles of tech companies such as Facebook and Google.
- We will investigate ethical issues in machine intelligence, such as privacy, surveillance, algorithmic fairness, elections, and international affairs.
- We will discuss how machine intelligence is portrayed in the media and in popular culture, and the impact of these portrayals on public views of science, AI, and technology.

Learning objectives

- This course will give you training in reading primary research articles, review papers, and book chapters - the majority of what scientists read. These activities are a different experience than reading textbooks. Reading and critical thinking are key skills in understanding science.
- This course will enable you to engage in constructive scientific conversations and debates, which will broaden and deepen your understanding of research in machine intelligence. We will discuss the sorts of questions you should ask about new results, and how to interpret them in the context of other studies.
- This course will improve how you communicate in oral presentations and written work.
- You will learn how to review the literature and find relevant peer-reviewed papers, allowing you to keep up to date in any field of science. You will understand the links between different areas of machine intelligence and become familiar with the development of this research over several decades. You will appreciate the challenges to this research, such as technological impediments and "AI winters", and learn about the latest advancements.

Role in the psychology curriculum

PSYC GU4236 is an advanced seminar, designed particularly for undergraduates who are majoring in Psychology or in Neuroscience and Behavior, for students participating in the Post-bac Psychology Program, and for Psychology Graduate Students. Students with a background in the computational sciences, engineering, and philosophy are also welcome to apply.

- For students pursuing the Psychology major or the Post-bac Certificate Program in Psychology, this course can be used to fulfill the advanced seminar requirement and/or the Group I Cognition & Perception distribution requirement.
- For students pursuing the Neuroscience & Behavior major, this course can be used to fulfill the P5 Advanced Seminar in Psychology requirement.
- Graduate students seeking to use this course towards M.A. or M.Phil. requirements must obtain prior permission from the Director of Graduate Studies.

Readings: There is no textbook required for this course

Readings will comprise scientific articles, book chapters, literature reviews, and commentaries in the fields of machine intelligence; videos and podcasts may be assigned. These resources will be available through Canvas.

Course requirements

1. **Class preparation and participation**. The assigned readings are designed to expand your knowledge on machine intelligence and to hone your critical thinking skills. The topics discussed during the seminar sessions are complex, leaving plenty of space to discuss and debate. Strong preparation and participation will enable us to have high-level and thought-provoking discussions.

2. **Reading reflections**. The day of each class period by 5 PM, you will be asked to submit a reading reflection on Canvas. Your reflections will help the professor prepare for class discussions. Each reflection should be 275-325 words, double spaced, with references using APA style format. Include at least one citation from the readings in a given week.

   The goals of reading reflections are to help students keep current on course topics and to help the instructor understand where students may have difficulty with the readings. The reading reflections will indicate which topics students are most intrigued by and, therefore, which areas may warrant more focus during class time. Since the goal of these assignments is to keep you up to speed and to guide my teaching and our class discussions, the assignments will be graded on a pass/fail basis. (I can only accept responses submitted before the deadline.)

   In your reflections, I encourage you to think about aspects of the readings that raise questions for you, or make you feel a certain way, or make you wonder about an issue that you have not thought about. The objective is to help you make sense of the reading. You are also encouraged to make connections between theory, research, and everyday experiences of AI by drawing on your life experiences. Reflections should not merely summarize the readings. The instructor will respond in writing to each of your reading reflections.

3. **Reading enables thoughtful discussion**. It is important to engage with the material during class discussions, since your active participation in these discussions will contribute to your final grade. If you feel that regularly contributing to class discussions is difficult for you, you should raise this issue with me in private as soon as possible. In such cases, we might be able to work out a way for you to participate thoughtfully through your reading responses or in other ways.

Generally speaking, effective class preparation and participation could include:

- Asking insightful or clarifying questions.
- Connecting the reading to other readings in the course or readings you've done on your own, drawing parallels and/or contrasts among ideas.
- Actively listening to fellow classmates and responding to their ideas.
- Offering thoughtful critiques of the research methodology and providing suggestions for how it might be improved.

- Bringing in outside sources – from the news media, blogs, podcasts, magazines, or the scientific literature – that shed light on machine intelligence findings or that illustrate ways in which these findings may be interpreted and applied.

4. **Present articles in class and comment on Discussion posts**. There is no shortage of media articles, blogs, links, etc. about machine intelligence. You will post four of them on the Discussion board on Canvas. Details to include in Discussion posts and presentations in class include: title of article, author, where it was published, date published, what it is about, how it relates to the course, and why you found it interesting. You will present two of your selections in class, for about 3-4 minutes each, followed by a brief class discussion. The instructor will email students before class to schedule presentations. Finally, you will comment on four Discussion posts by other students (at least 3-5 sentences). **You can still earn extra credit by posting an additional four topics on the Discussion board (see below).

5. **Research paper and in-class presentation**. The culmination of this course is the creation of a research paper relating to the material of the class. Good writing is good thinking, and a primary goal of this assignment is to help you hone your writing and critical thinking skills. The process of writing the research paper follows three steps:

   First, you will be asked to identify a topic related to the class that you may want to write about and present to your fellow students. You should email the instructor stating your research topic, so that together we can decide whether it is appropriate, and fine-tune it. Topic proposals should include a short paragraph about your intended topic and a list of at least five references you intend to use. I will make suggestions regarding the content, resources, etc.

   Second, once your topic is approved, you will begin to work on a first draft of the paper. Generally, you will want to choose a topic that is appropriately focused for an 8 page paper (minimum 8 pages: double spaced, 11 point font, where the references do not count towards the eight pages; you may include up to 2 figures, which do not count towards the eight pages). The paper may include the following sections: Introduction to the topic; Background in which you review knowledge and advances in the field; Sections on specific sub-topics on your project; a Conclusion or Summary section to present your ideas, analysis, and future directions; and a References section.

   Third, I will provide comments and suggestions on your first draft, and you will be expected to make substantive changes – not just copy editing, but larger edits, such as reworking entire sections, drawing on new sources, and providing more analysis. The final draft of the paper will be graded not only as a standalone paper, but also on how it demonstrates improvement over the earlier draft.

Grading

Grades will be calculated based on the percentages outlined below.

- **Class preparation and participation**: 25%
  - Reading reflections: 15%
  - Contribution to class discussion: 10%
- **Discussion**: Present AI articles in class, post articles onto Canvas Discussion section, and comment on Discussion posts by other students: 30%
  - Each student will discuss in class (4 minutes, followed by brief discussion) two examples of news articles, research articles, or blog posts about machine intelligence.
    - Post your articles on the Discussion board on Canvas. I will contact students before each class about who will present. Details to include in your Discussion post and video presentation: title of article, author, where it was published, date published, what it is about, how it relates to the course, and why you found it interesting.
  - Each student will post four relevant articles in the Discussion section on Canvas.
  - Each student will provide brief comments (about 3-5 sentences) on four Discussion posts on Canvas by other students.
- **Main Project**: Research paper and in-class presentation: 45%
  - Choose topic, in consultation with instructor, by Wednesday Dec. 9, end of day.
  - In class presentation of project (20% of grade): Wednesday Dec. 16
  - First draft of written paper (5% of grade): Monday Dec. 16 by 11:59 pm EST.
  - Final version of written paper (20% of grade): Wednesday Dec. 23 by 6 pm EST.
- **Extra credit**
  - You may earn extra credit if you post an additional four links on the Discussion board.

Module 1: September 9, 2020

Introduction to the seminar

Topics / Questions:

- What is machine intelligence and artificial intelligence?
- What are examples of machine intelligence?
- How does AI impact you?
- Are you excited about AI, concerned about it, or both?
- What topics in AI interest you?

Module 2: September 16, 2020

Foundations of machine intelligence

Topics / Questions:

- What is intelligence?
- Can machines be intelligent?
- What kinds of machines are used in machine intelligence?
- How can we assess whether machines are intelligent?

Readings

Turing, A. M. (1948). Intelligent machinery. Chapter 10 in *The Essential Turing*, ed. J. Copeland. Clarendon, 2004.

Turing, A. M. (1950). Computing machinery and intelligence. *Mind* 59 (236): 433–460.

Samuel, A. L. (1959). Some studies in machine learning using the game of checkers. *IBM Journal of research and development*, 3(3), 210-229.

Supplemental reading:

Lefkowitz, M. (2019). Professor's perceptron paved the way for AI – 60 years too soon. *Cornell Chronicle*, September 25, 2019.

Module 3: September 23, 2020

Algorithms, data, compute, and technology

Topics / Questions:

- Discuss elements of machine intelligence:
  - **data** used in an AI model;
  - **algorithms**, i.e., the sets of rules or instructions given to an AI program;
  - **compute** (computing resources);
  - **technology**, such as a camera for capturing images, a wearable sensor for measuring hand tremor, or a social media platform, all of which generate data.
- Discuss AI systems and identify the roles of data, algorithms, compute, and technology

Readings

Marr, D. (1982). Section on "The Three Levels" in Chapter 1 of *Vision* by David Marr.

Gershgorn, D. (2017). The data that transformed AI research - and possibly the world. *Quartz*, July 26, 2017.

Duke University PULSE study.

Cai, F. (2020). Yann LeCun Quits Twitter Amid Acrimonious Exchanges on AI Bias. *Synced*, June 30, 2020.

Module 4: September 30, 2020

How machines acquire knowledge and skills

Topics / Questions:

- What is "learning" in machine intelligence?
- How do machines acquire "knowledge" or "abilities"?
- Data-driven learning vs. built-in knowledge or priors
- Nature and nurture
- Rationalism and empiricism

Readings

Katz, Y. (2012), Noam Chomsky on where artificial intelligence went wrong. *The Atlantic*.

Gopnik, A. (2019). Will A.I. ever be smarter than a four-year-old? *Smithsonian Magazine*. (Watch the videos)

Heaven, D. (2019). Deep trouble for deep learning. *Nature* 574: 163-166.

Waldrop, M. M. (2019). What are the limits of deep learning? *PNAS* 116 (4): 1074-1077.

Module 5: October 7, 2020

Neural networks and machine learning

Guest speaker: Dr. James Hedges, Principal Data Scientist, T Rowe Price; and Visiting Scientist, Zuckerman Mind Brain Behavior Institute, Columbia University

Topics / Questions:

- Introduction to machine learning
- History of neural networks and biological inspiration in brain science
- Basics of neural networks

Readings

Mitchell, M. (2019). Neural networks and the ascent of machine learning. Ch. 2 in *Artificial Intelligence: A Guide for Thinking Humans* by Melanie Mitchell.

Mitchell, M. (2019). Who, what, when, where, why. Ch. 4 in *Artificial Intelligence: A Guide for Thinking Humans* by Melanie Mitchell.

Module 6: October 14, 2020

Deep learning

Guest speaker: Dr. Nikolaus Kriegeskorte, Professor, Visual Inference Lab, Zuckerman Mind Brain Behavior Institute, Columbia University

Topics / Questions:

- Introduction to deep learning

Reading

Kriegeskorte, N. (2015). Deep neural networks: A new framework for modeling biological vision and brain information processing. *Annual Review of Vision Science* 1: 417-446.

Module 7: October 21, 2020

Reinforcement learning

Guest speaker: Dr. Bruno B. Averbeck, Chief, Section on Learning and Decision Making, National Institute of Mental Health, NIH

Topics / Questions:

- Origins of reinforcement learning in psychology and neuroscience
- Introduction to reinforcement learning
- Applications of reinforcement learning in artificial systems

Readings and materials

Watch the documentary AlphaGo before class.

Neftci, E. O., Averbeck, B. B. (2019). Reinforcement learning in artificial and biological systems. *Nature Machine Intelligence* 1 (3), 133-143.

Supplemental reading:

Koch, C. (2016). How the computer beat the Go player. *Scientific American* 27 (4), 20-23.

Module 8: October 28, 2020

Cognitive processes in machine intelligence

Topics / Questions:

- How are cognitive processes embedded in machines?
- Discuss ideas such as *intuitive psychology* and *intuitive physics*

Readings

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences* 40, E253.

Module 9: November 4, 2020

Reverse engineering the mind

Guest speaker: Dr. Josh McDermott, Associate Professor, Laboratory for Computational Audition, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology

Topics / Questions:

- Approaching cognitive science and neuroscience like an engineer who wants to build an intelligent system

Readings and video

Franci, A., & McDermott J. H. (2020). Deep neural network models of sound localization reveal how perception is adapted to real-world environments. bioRxiv.

Interview with Josh Tenenbaum in *Architects of Intelligence* (2018), ed. Martin Ford.

Tenenbaum, J. (2020). Technology Day 2020: Geniuses and Game Changers - Josh Tenenbaum.

Module 10: November 11, 2020

Robotics

Topics / Questions:

- Introduction to robotics and robot learning
- Concept of physical artificial intelligence, as distinguished from data-drive AI
- Areas within robotics, such as soft robotics, aerial and aquatic robotics, dexterous movements, locomotion, and autonomous vehicles

Readings

Kaelbling, L. P. (2020). The foundation of efficient robot learning: Innate structure reduces data requirements and improves robustness. *Science* 369 (6506), 915-916.

Man, K, & Damasio, A. R. (2019). Homeostasis and soft robotics in the design of feeling machines. *Nature Machine Intelligence* 1, 446–452.

Hsu, J. (2019). Machines on mission possible. *Nature Machine Intelligence* 1, 124–127.

Module 11: November 18, 2020

Human-machine interaction; meta-cognition and consciousness in artificial systems

Guest speaker: Dr. Megan Peters, Assistant Professor, Cognitive & Neural Computation Lab, University of California, Irvine.

Topics / Questions:

- Interactions between people and machines
- The tendency of humans to attribute higher level cognition to machines
- Meta-cognition and consciousness in artificial systems

Readings

Metcalfe, J. (2013). Anoetic, noetic, and autonoetic metacognition. In *Foundations of Metcognition*, eds. Michael J. Beran, Johannes Brandl, Josef Perner, and Joëlle Proust.

Thellman, S., Silvervarg, A., Ziemke, T (2017). Folk-psychological interpretation of human vs. humanoid robot behavior: Exploring the intentional stance toward robots. *Frontiers in Psychology* Volume 9, article 1962.

Module 12: December 2, 2020

Brain-machine interface

Guest speaker: Dr. Thomas Naselaris, University of Minnesota. Dr. Naselaris is an expert on brain decoding and did research in Dr. Jack Gallant's lab, which is mentioned in the reading.

Topics / Questions:

- Decoding brain signals to control machines
- Neural prosthetics
- Incorporating machinery into humans, e.g., deep brain stimulation

Readings

Velasquez-Manoff, M. (2020). The Brain Implants That Could Change Humanity. *New York Times*, August 28, 2020.

Module 13: December 9, 2020

Facial recognition

Topics / Questions:

- Technology of facial recognition
- Ethics of facial recognition

Readings

Brown, L. (2019). There will be no turning back on facial recognition. *New York Magazine* Nov. 12, 2019.

Roussi, A. (2020). Resisting the rise of facial recognition. *Nature* 587: 350-353.

Castelvecchi, D. (2020). Beating biometric bias. *Nature* 587: 347-349.

Module 14: December 16, 2020

In class presentations by students on their research projects

Reading

Jordan, M. I. (2019). Artificial intelligence - the revolution hasn't happened yet. *Harvard Data Science Review* 1.1.