**** The topics and readings in the syllabus will change during the semester. Please consult the syllabus weekly**

Department of Psychology – Columbia University

Machine Intelligence

GU4236 - Spring 2022

4 points

Instructor: Trenton Jerde, Ph.D.

Class Meets: M 4:10 pm - 6:00 pm, 405 Schermerhorn Hall [SCH]

Email: taj2128@columbia.edu

Office Hours: Online, by appointment

Instructor
The instructor is a Senior Editor at *Nature Machine Intelligence*, a scientific journal on artificial intelligence, machine learning, and robotics. He has a BA in psychology from the University of Iowa, a PhD in neuroscience from the University of Minnesota, and postdoctoral training in cognitive neuroscience from New York University. He has held faculty positions at Columbia University, New York University, and the University of Minnesota.

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.

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.
- We will investigate ethical issues in machine intelligence, such as privacy, surveillance, algorithmic fairness, 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.

Module 1: January 24, 2022

Introduction to the seminar

Topics / Questions:

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

Module 2: January 31, 2022

Foundations and perspectives on AI

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

Strickland E (2021). The turbulent past and uncertain future of AI. *IEEE Spectrum*, October 2021.

Jordan MI (2019). Artificial intelligence – The revolution hasn't happened yet. *Harvard Data Science Review* 1.1.

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

Module 3: February 7, 2022

History and concepts of AI

Topics / Questions:

- Appreciate the personal and professional narratives of some researchers
- Learn about the evolution of certain core ideas, and how researchers navigated the work culture

- Get a big-picture, historical, and multidisciplinary overview of AI

Readings

Metz C (2021). Promise. Chapter 2 in *Genius Makers: The Mavericks Who Brought AI to Google, Facebook, and the World*, by Cade Metz. Dutton.

Metz C (2021). Rejection. Chapter 3 in *Genius Makers: The Mavericks Who Brought AI to Google, Facebook, and the World*, by Cade Metz. Dutton.

Russell S, Norvig P (2021). Introduction. Chapter 1 in *Artificial Intelligence, a Modern Approach*, 4th Edition, by Stuart Russell and Peter Norvig. Pearson.

Module 4: February 14, 2022

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

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

Editorial (2022). The rise and fall (and rise) of datasets. *Nature Machine Intelligence* 4, 1–2.

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

Ethical and societal considerations:

Duke University PULSE study (skim this)

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

Module 5: February 21, 2022

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

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

Sutton R (2019). The bitter lesson.

Module 6: February 28, 2022

Recent trends on data, algorithms, compute, and learning

Topics / Questions:

- The shift in emphasis towards data curation
- The nature of compute across the decades
- The merging of data, algorithms, and technology to address scientific questions

Data-centric AI Resource Hub. Watch the two-minute video by Andrew Ng.

Sevilla et al. (2022). Compute trends across three eras of machine learning. arXiv.

Preston E (2020). Baby headcams reveal the world that infants actually see.

Orhan AE, Gupta VV, Lake BW (2020). Self-supervised learning through the eyes of a child. *Advances in Neural Information Processing Systems* 33. (Read the Abstract, Introduction, Dataset section, Limitations, and Discussion; no need to focus on the technical details.)

Module 7: March 7, 2022

Neural networks and deep learning

Topics / Questions:

- Neural networks and their biological inspiration in brain science
- Introduction to deep learning

Readings

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

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

Moore SK, Schneider D, Strickland E (2021). How deep learning works. *IEEE Spectrum*, October 2021, 32-33.

Kriegeskorte N. Introduction to deep neural networks (video in the Files section on Courseworks).

Background reading (optional):

Sejnowski TJ (2020). The unreasonable effectiveness of deep learning in artificial intelligence. *Proceedings of the National Academy of Sciences* 117 (48).

Module 8: March 21, 2022

AI movie review!

Coded Bias: YouTube link

AlphaGo: YouTube link

Module 9: March 28, 2022

Limitations of deep learning and possible future directions of AI

Topics / Questions:

- What is missing from current deep learning, and why?
- What role will deep learning play in AI going forward?
- What other approaches may be incorporated in machine intelligence along with deep learning?

Readings

Marcus G (2022). Deep Learning Is Hitting a Wall. *Nautilus*, March 10, 2022.

Machine Learning Street Talk, episode #70, interview with Letitia Parcalabescu (2022). YouTube link. Listen to the introduction; the conversation with Letitia is optional (and interesting).

Pearle J (2019). The limitations of opaque learning machines. Chapter 2 in *Possible Minds: 25 Ways of Looking at AI*. Edited by John Brockman. Penguin Press.

Geirhos R, Jacobsen JH, Michaelis C et al. (2020). Shortcut learning in deep neural networks. *Nature Machine Intelligence* 2, 665–673.

Module 10: April 4, 2022

Cognitive approaches to machine intelligence

Topics / Questions:

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

Readings

Tenenbaum J (2022). What kind of computation is cognition? The spring 2022 Shulman Lectures at Yale. Watch the first 30 minutes, until the slide on How do we build this architecture; the rest is optional. Link

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

Tenenbaum JB, Kemp C, Griffiths TL, Goodman ND. (2011). How to grow a mind: statistics, structure, and abstraction. *Science* 331(6022):1279-85.

Module 11: April 11, 2022

Neuroscience approaches to machine intelligence

Topics / Questions:

- The relationship between neuroscience and artificial intelligence
- What levels of neuroscience, and what kinds of information, are potentially useful for AI systems?

Readings

Ullman S (2019). Using neuroscience to develop artificial intelligence. *Science* 363 (6428): 692-693.

Zador AM (2019). A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications* 10:3770.

Brain Inspired Podcast 067, Paul Cisek: Backward Through The Brain.

Module 12: April 18, 2022

Robotics

Topics / Questions:

- Introduction to robotics and robot learning
- The concept of physical artificial intelligence, as distinguished from data-driven AI

Readings

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

Editorial (2020). Materializing artificial intelligence. *Nature Machine Intelligence* 2, 653.

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

Additional reading, which is optional, but the first page is relevant to our previous discussions about learning vs. innate structures:

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

Module 13: April 25, 2022

Brain-machine interface and human-computer interaction

Topics / Questions:

- Decoding brain signals to control machines, e.g., neural prosthetics
- Incorporating machinery into humans, e.g., deep brain stimulation
- Human and machine interactions

Readings

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

Khatchadourian R (2018). Degrees of freedom. *The New Yorker* Vol. XCIV 38: 56.

Heider and Simmel animation (1944). YouTube Link

Before Siri and Alexa, there was ELIZA. YouTube Link. Wikipedia description

Mother meets her deceased daughter through VR technology. YouTube Link. Article about the story (link)

Module 14: May 2, 2022

Autonomous vehicles, artificial writing, and class summary

- Artificial writing: in what ways can machines "write"?
- Autonomous vehicles: history, challenges, and current status
- Review of course topics and key points

Hutson M (2021). The language machines. *Nature* 591: 22-25.

Johnson S (2022). A.I. is mastering language. Should we trust what it says? *New York Times*, April 15, 2022.

Bender E (2022). On NYT Magazine on AI: Resist the urge to be impressed. *Medium*, April 17, 2022.

Lipson H, Kurman M (2016). The robotic chauffeur. Chapter 1 in Driverless: Intelligent Cars and the Road Ahead. MIT Press.

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.

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 Courseworks.

Grading

Grades will be calculated based on the percentages outlined below.

- **Class preparation and participation**: 30%
    - Reading reflections: 15%
    - Contribution to class discussion: 15%
- **Discussion**: Post at least four articles onto the CourseWorks Discussion section, and comment on at least four Discussion posts by other students: 20%

    - Details to include in your Discussion post: web link, 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.
- **Main Project**: Research paper: 50%
    - Choose topic, in consultation with instructor, by **April 11**, 2022 by 11:59 pm EST.
    - First draft of written paper (10% of grade): **April 25**, 2022 by 11:59 pm EST.
    - Final version of written paper (40% of grade): **May 13** by 3:00 pm EST.

Course requirements

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.

Reading reflections. The day of each class period by 3 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.

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.

Post comments or links to journalistic articles or videos or science papers or blogs etc. in the Discussion forum and comment on Discussion posts of others. 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. Additionally, you will **comment on four** Discussion posts by other students (at least 3-5 sentences).

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.