Philosophy 340: Minds and Machines (Fall 2023)
Professor Eric Swanson

This course will discuss minds, machines, the relationships between them, and the relationships they encourage and discourage. Throughout we will compare and contrast human minds and brains with machine 'minds' and 'brains,' considering questions like: Could a machine have a mind? Could a machine be conscious, or think in the ways that people do? What can human brains teach us about how a thinking machine might work, and vice versa? What might information processing, information integration, and information availability have to do with consciousness? How do machines and our interactions with them influence how humans think, learn, reason, and know? What are the promises and perils of artificial intelligence, big data, and gargantuan networks, especially when the underlying mechanics of machine and network reasoning are not transparent to us? How should machines be treated—including machines that exhibit some but not all features of human minds? Why? How should machines treat us?

Course web sites

Canvas for lecture, Canvas for sections

Keeping in touch

Use Piazza (on Canvas) to ask questions that aren't specific to you! This helps everyone, especially because in this course everyone has different strengths and areas of expertise.

Use a private post on Piazza, addressed to both me and your GSI, for questions specific to you.

My open office hours are 9:15–11:15 on Tuesdays, over Zoom (https://zoom.us/my/e.swanson). You don't need to make an appointment. I'm also happy to meet at other times by appointment.

Books you should buy or borrow

A Brief History of AI, by Michael Wooldridge

Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy, by Cathy O'Neil

These are available as Kindle editions. Other texts and media are available on Canvas.

Lecture

Mondays and Wednesdays 10-11:20, 2500 SKB. Please attend lecture in person if you're healthy and able. You can also watch on Lecture Capture, but be sure to keep up!

Discussion Sections

002: Fridays 9–10 with Sarah Hintzen, in 2070 SKB

003: Fridays 1–2 with Mica Rapstine, in 2271 AH

004: Fridays 2–3 with Mica Rapstine, in 2271 AH

005: Fridays 10–11 with Sarah Hintzen, in 3401 MH

006: Thursdays 3–4 with Christopher Feng, in G311 DENT

007: Thursdays 4–5 with Christopher Feng, in 1508 NUB

008: Thursdays 2–3 with Humaira Nasir, in 2234 USB

009: Fridays 1–2 with Humaira Nasir, in 3463 MH

010: Wednesdays 2–3 with Justin Peterson, in 1636 CHEM

011: Wednesdays 3–4 with Justin Peterson, in 3752 SSWB

Grading

Submit assignments on the Canvas site for your section.

2%: self-introduction, due on Canvas by 9 pm on 6 Sep. (or a week after you add the course), graded pass / fail.

1% each (13 total): a discussion question for section, due on Canvas when your GSI decides. These responses should clearly articulate a question you'd be interested in discussing in section, with enough background supplied so it's clear why you have that question and how it relates to a contemporaneous reading or lecture. Graded pass / fail.

20% each: two three to four page papers, due by 9 pm on 25 Sep. and 13 Nov.

10%: midterm exam, in lecture, 23 Oct.

20%: final exam, 10:30–12:30, 13 Dec.

15%: your section grade—participation in section and (optionally) in lecture and on Piazza, comprehension quizzes if needed, etc. Excellent participation in lecture and on Piazza will boost your section grade.

Expectations

Things may change, so pay attention to announcements in lecture, in section, and on Canvas.

Keep up with the readings, and do them *before* lecture. They demand careful, patient thought and attention. If you fall behind you'll miss much of what's going on for the rest of the semester.

We'll be discussing controversial issues in this course, and challenging our beliefs throughout. Be polite, respectful, and generous in discussion with your classmates and teachers.

We expect all students to understand the College of LSA's standards of academic integrity:

https://lsa.umich.edu/lsa/academics/academic-integrity.html
https://www.lib.umich.edu/academic-integrity

Contact your GSI if you don't understand them fully! We encourage you to discuss the material with your classmates, but all work you turn in must be your own. Cheating or plagiarism will result in a grade of 0 for that assignment and disciplinary action by the University.

If you might need an accommodation for a disability, let us know right away. Some aspects of this course may be modified to facilitate your participation and progress. The Office of Services for Students with Disabilities (734-763-3000; ssd.umich.edu) typically recommends accommodation through a Verified Individualized Services and Accommodations form. Any information you provide is private and confidential.

Use of laptops in class can distract the user, other students, and teachers. Lots of research has shown that many students learn better if they take notes by hand. See, for example: http://www.scientificamerican.com/article/a-learning-secret-don-t-take-notes-with-a-laptop/. However, some students absolutely have to use laptops, and do benefit from them. If you need to use a laptop in lecture, please sit to the left (as you face the front of the room). If laptops distract you, please sit in the center or on the right. No phones, texting, etc. in class.

Missed exams can be made up but only with appropriate documentation—for medical reasons, a signed and dated doctor's note; for the death of a loved one, a statement of death; etc. You must contact your GSI in advance to ask for an extension on a paper due date.

So you can all pay attention to all of lecture, we will finish by 11:18. But I ask you to be fully present for our discussion until 11:18.

University of Michigan is committed to advancing the mental health and wellbeing of its students. If you or someone you know is feeling overwhelmed, depressed, and/or in need of support, services are available. For help, contact Counseling and Psychological Services (CAPS) at (734) 764-8312 and https://caps.umich.edu/ any time, or through its counselors physically located in schools on both North and Central Campus. You may also consult University Health Service (UHS) at (734) 764-8320 and https://www.uhs.umich.edu/mentalhealthsvcs or for alcohol or drug concerns, see http://www.uhs.umich.edu/aodresources. For a listing of other mental health resources available on and off campus, visit: http://umich.edu/~mhealth/.

Mon, 28 Aug.	**Organizational meeting**
Wed, 30 Aug.	**Pre-history; computability; representation; levels of explanation** Michael Wooldridge, *A Brief History of Artificial Intelligence,* Introduction and Chapter 1 Luiz Pessoa, *The Entangled Brain,* Introduction and Chapter 1
Wed, 6 Sep.	**Good old fashioned artificial intelligence (GOFAI)** Michael Wooldridge, *A Brief History of Artificial Intelligence,* Chapter 2 Jill Lepore, "All the King's Data" aka "How the Simulmatics Corporation Invented the Future" ⇒ **Self-introduction due on Canvas by 9 pm (or a week after you add the course)**
Mon, 11 Sep.	**Knowledge, expertise, and intelligence** Michael Wooldridge, *A Brief History of Artificial Intelligence,* Chapter 3 Hubert Dreyfus and Stuart Dreyfus, "Peripheral Vision: Expertise in Real World Contexts"
Wed, 13 Sep.	**Robotics, embodiment, enaction** Michael Wooldridge, *A Brief History of Artificial Intelligence,* Chapter 4 Arthur Glenberg, Jessica Witt, and Janet Metcalfe, "From the Revolution to Embodiment: 25 Years of Cognitive Psychology"
Mon, 18 Sep.	**Using a possible future to help us reflect on the present** *Her,* written and directed by Spike Jonze
Wed, 20 Sep.	**Deep learning, part 1** Michael Wooldridge, *A Brief History of Artificial Intelligence,* Chapter 5 Rodney Brooks, "Machine Learning Explained"
Mon, 25 Sep.	**Deep learning, part 2** Michael Wooldridge, *A Brief History of Artificial Intelligence,* Chapter 6 Sara Hooker, "The Hardware Lottery" Rich Sutton, "The Bitter Lesson" ⇒ **First paper due by 9 pm**
Wed, 27 Sep.	**Brains, part 1** Francois Jacob, "Evolution and Tinkering" Luiz Pessoa, *The Entangled Brain,* Chapter 2
Mon, 2 Oct.	**Brains, part 2** Luiz Pessoa, *The Entangled Brain,* Chapters 3–4
Wed, 4 Oct.	**Emotion, part 1** Luiz Pessoa, *The Entangled Brain,* Chapter 5
Mon, 9 Oct.	**Emotion, part 2** Luiz Pessoa, *The Entangled Brain,* Chapter 6 Lisa Feldman Barrett, "The Big Idea: Do Animals Have Emotions?"

Wed, 11 Oct.	**Natural and artificial attention** Luiz Pessoa, *The Entangled Brain,* Chapter 7 Stephen Wolfram et al., excerpt from "What Is ChatGPT Doing . . . and Why Does It Work?" (stop reading right before the section "Meaning Space and Semantic Laws of Motion")
Wed, 18 Oct.	**Some problems with foundation models** Alison Gopnik, "Large Language Models as a Cultural Technology" Rodney Brooks, "What Will Transformers Transform?" Ted Chiang, "Catching Crumbs from the Table"
Mon, 23 Oct.	⇒ **Midterm exam**
Wed, 25 Oct.	**Complexity, AI, college admissions, and online advertising** Cathy O'Neil, *Weapons of Math Destruction,* Chapters 1, 3, and 4 Luiz Pessoa, *The Entangled Brain,* Chapter 8
Mon, 30 Oct.	**AI, justice, and employment** Cathy O'Neil, *Weapons of Math Destruction,* Chapters 5–7
Wed, 1 Nov.	**AI, politics, and law** Cathy O'Neil, *Weapons of Math Destruction,* Chapter 10 and Conclusion Gary Marcus, "Automatic Disinformation Threatens Democracy—and It's Here" Gary Marcus, "The Imminent Enshittification of the Internet"
Mon, 6 Nov.	**Evolution and integration** Luiz Pessoa, *The Entangled Brain,* Chapters 9 and 10
Wed, 8 Nov.	**Unlearning fear** Luiz Pessoa, *The Entangled Brain,* Chapter 11 Golnaz Tabibnia and Dan Radecki, "Resilience Training that can Change the Brain"
Mon, 13 Nov.	**Looking back** Luiz Pessoa, *The Entangled Brain,* Chapter 12 ⇒ **Second paper due by 9 pm**
Wed, 15 Nov.	**Sci-fi futures** Michael Wooldridge, *A Brief History of Artificial Intelligence,* Chapter 7 Rodney Brooks, "The Seven Deadly Sins of Predicting the Future of AI"
Mon, 20 Nov.	**AI, alienation, bias, fakes, and glitches** Michael Wooldridge, *A Brief History of Artificial Intelligence,* Chapter 8 Ruha Benjamin, "Default Discrimination"
Mon, 27 Nov.	**Ideology, human intelligence, and artificial intelligence** *The Matrix,* written and directed by Lilly and Lana Wachowski
Wed, 29 Nov.	**Consciousness and what it's like** Michael Wooldridge, *A Brief History of Artificial Intelligence,* Chapter 9 Rodney Brooks, "What Is It Like to Be a Robot?"

Mon, 4 Dec.	**Prediction, directions of fit, and what we value** "Hang the DJ" (*Black Mirror* Season 4) Abeba Birhane, "The Impossibility of Automating Ambiguity"
Wed, 6 Dec.	**Looking back, looking forward; exam review**
Wed, 13 Dec.	⇒ **Final exam from 10:30–12:30, location TBD.**