

Machine Intelligence

Department of Psychology – Columbia University
Psych # – Course Syllabus – Spring Semester 2026
Class Days TBD
Location TBD

Instructor: Trenton Jerde, Ph.D.
Email: taj2128@columbia.edu
Office Hours: Online, by appointment

Teaching Assistant:
Email:
Office Location:
Office Hours:

Bulletin description. This course 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, machine learning, robotics, connectionism, neural networks, and deep learning. Emphasis will be placed on the conceptual understanding of topics. The course does not include, nor require a background in, computer programming and statistics. The overall goal is for students to become informed consumers of artificial intelligence applications.

Course description: In this course, students will trace the history of artificial intelligence, learning about the key ideas, motivations, and debates that have shaped the field. We will ask fundamental questions: What is intelligence? Can machines think? Should the goal of machine intelligence be to emulate human cognition – even extending toward artificial general intelligence (AGI) or even superintelligence – or instead to remain a set of tools designed to assist human life? Along the way, students will learn to distinguish the components of machine intelligence—data, algorithms, compute, and technology—and to see how these elements combine in applications that range from chatbots to drones to robots.

The course will introduce major approaches to machine learning, artificial neural networks, and deep learning, while also examining how artificial systems compare with biological ones in learning and knowledge acquisition. We will consider how core psychological and neural processes such as perception, language, thought, emotion, and movement can be modeled in machines. Students will also explore real-world applications of AI, from facial recognition to large-scale deployment by tech companies, and will critically evaluate the ethical concerns these raise, including issues of privacy, surveillance, and algorithmic fairness. Finally, we will discuss how machine intelligence is represented in the media and popular culture, and how these portrayals shape public understanding of AI, science, and technology.

Course Readings: There is no required textbook for this course. Readings will include scientific articles, book chapters, literature reviews, and commentaries related to machine intelligence. Additionally, videos and podcasts will be incorporated. All materials will be accessible through Courseworks.

Prerequisites: 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.

Role in the curriculum: PSYC TBD is an undergraduate, intermediate-level lecture course, open to undergraduates and students in the post-baccalaureate Psychology program. It can be used to fulfill the following requirements:

- For the Psychology Major, this course will fulfill the Group 1 requirement or the Special Elective Requirement
- For the Postbac Certificate Program, this course will fulfill the Group 1 requirement.
- For the Neuroscience and Behavior Major, this course will fulfill the P4 requirement.

Course website: The course website (on Courseworks) contains the most up to date information. This syllabus is subject to change, so make sure to check the course website for the most current version, as well as announcements for changes in the schedule.

Important dates

- **Midterm exam #1:** TBD
- **Midterm exam #2:** TBD
- **Exam #3 (Final Exam):** TBD

Attendance and preparation

Lectures: You will not do well in the class if you don't attend lecture. Lectures may or may not involve reiteration or expansion of material in the readings and materials (videos, podcasts). Nevertheless, you are still responsible for reading and understanding each chapter as assigned. Feel free to ask questions during class or office hours if there are difficulties with material in the readings. To ensure that you and the other students get the benefits of this class, you should complete the assigned readings prior to class.

Classroom Decorum. To maintain a classroom environment that is both respectful of others and conducive to learning we ask that you observe the following:

- Lectures start promptly at time TBD, and we ask you to be in your seat by that time. Please do not disrupt the class by coming late or leaving early.

- Turn off your cell phone during class.
- You may use a laptop or other device for note taking. Do not use your electronic devices for non-course-related activities. Doing so is not only distracting to those around you but also broadcasts disrespect for the class.

Evaluation:

1. (40%) Highest Midterm Exam Grade (out of 2)
2. (50%) Final Exam
3. (10%) In-class Writing Assignments
4. (1 point) Extra Credit (from Midterm route)
5. (1 point) Extra Credit (in-class writing assignment route)

Course Score	Course Grade
97-100	A+
93-96	A
90-92	A-
87-89	B+
83-86	B
80-82	B-
77-79	C+
73-76	C
70-72	C-
60-69	D
0-59	F

1. Midterm examinations: Two midterms occur during the semester. Your **lower score out of the two midterm tests will be dropped**. Although exams will tend to focus more on the material covered in lecture, any and all material in the lectures and texts is fair game.

The teaching assistants will keep the exams and answer sheets. If you want to see your exam again or look at it for an extended period of time (and I strongly encourage you to do so), then you will want to visit one of the teaching assistants during office hours or make an appointment.

About showing up late for an exam: Do not be late for an exam. If you arrive after someone else has already finished the test and left the room, you will not be allowed to take the exam.

Exam Makeups. There are **no provisions for making up exams**, because the opportunity to miss one midterm with no penalty is built into the grading system. Please take a moment **now** to note the dates of exams and in-class review sessions. Because of the size of this class, it will not be possible to schedule alternative exam times for students who have difficult examination

schedules, travel plans, etc. An unexcused absence from any of the exams will result in a grade of F on that exam.

Midterm Assessments. We will have two midterms, which will each include a mix of short-answer and multiple-choice questions. We will post some sample questions and a review sheet on Canvas before each exam. Midterms are designed to be completed in 60 minutes, though you will have the full 75-minute class period to complete them.

The higher of your two midterm scores will count toward your final grade; the lower score will be dropped. That means that you may skip one of the midterms, though I strongly recommend that you plan to take both, since testing yourself on course material helps to build long-term learning, and is one of the best ways to begin studying for the final. Midterms will be held in class on the following dates:

Midterm 1: TBD

Midterm 2: TBD

There are no make-up midterms offered in this course. If you need to miss a midterm for any reason (excused or unexcused), that will be the score that is dropped. You are still encouraged to look over a skipped midterm as preparation for the next exam(s)—you may do this in any of our office hours as long as you let us know ahead of time so we can have a blank exam on hand. If you know that you will have a conflict with both midterm dates, you should not take this class. If you are a student athlete and anticipate that you might be traveling during one or both of these dates, please come talk to me before the end of the second week of class to see if it will be feasible for you to take your midterms on the road.

Because your lower midterm score is dropped, and because some exam questions will be closely related to material that was asked on problem sets (and answered in the answer keys), the midterms for this course are geared to be slightly more challenging than those in some other PSYC courses. Questions will focus more on applying, connecting, and analyzing course concepts, rather than defining them. If you don't do as well on the first exam as you would have liked, don't despair—that score will go away entirely if you do better on the second midterm.

2. Final examination: The final exam occurs during Finals week. Note – scheduling is registrar controlled, and it is currently scheduled for time TBD. Unfortunately, I cannot control the time of the final. Therefore, I advise you now to consider your spring travel schedule. Students will not be able to reschedule the exam because of travel needs. Although the final will tend to focus more on the material covered in lecture, any and all material in the lectures and weekly study materials is fair game. The final will be cumulative, and you must take the Final Exam, although it will focus more heavily on the material from the last 1/3 of the course.

Final Exam. The final will cover the material from the whole course, with more emphasis on topics from after Midterm 2. It will have the same structure as the midterms, but slightly longer: it will be designed to be completed in about 90 minutes, though you'll have the full 170-minute exam period to complete it.

Projected Final Exam Date: In-person, TBD

If you have a conflict with this final exam date, you should speak to the instructor before enrolling in this class. Although the Registrar could shift exams from the originally projected date, it is overwhelmingly likely that our final will take place on this date. Columbia University only allows final exams to be taken outside of the scheduled slot in extreme circumstances such as a medical or family emergency. Columbia’s exam policy is that “Examinations are not rescheduled to accommodate travel plans.” The policy of this course is that early exams are not offered for any reason.

3. Writing assignment: There will be three brief in-class writing assignments. For these writing assignments, you will need a pencil and eraser.

About missed writing assignments: Your lowest score out of the three writing assignments will automatically be dropped. Therefore, there are no provisions for making up the writing assignments, because the opportunity to miss one with no penalty is built into the grading system.

4. Extra Credit (Midterm route): If you take and pass both **midterm** exams, you will receive an extra point towards your final grade.

5. Extra Credit (in-class writing assignment route): If you turn in all 3 in-class writing assignments (note – you must be in class on these days), you will receive an extra point towards your final grade.

Tentative List of Course Topics

Class 1	Introduction to machine intelligence
Class 2	History of AI: Emergence from the fields of philosophy, mathematics, psychology, and neuroscience
Class 3	Foundations and perspectives on machine intelligence
Class 4	On intelligence, artificial general intelligence, and superintelligence
Class 5	Elements of an AI system: Data, algorithms, compute, and technology
Class 6	Learning: How machines acquire knowledge and skills
Class 7	Machine learning and artificial neural networks
Class 8	Language and AI, part 1: How LLMs work
Class 9	Language and AI, part 2: Meaning and reasoning in LLMs
Class 10	Midterm 1
Class 11	Emotion, motivation, and affective Computing
Class 12	AI companions and mental health
Class 13	Consciousness and AI
Class 14	Neuroscience and AI
Class 15	Computational neuroscience and representation
Class 16	Neural plasticity and lifelong learning
Class 17	Cognitive science and AI
Class 18	Developmental psychology and AI

Class 19	Social Cognition and Theory of Mind in Humans and Machines
Class 20	Midterm 2
Class 21	Embodied AI and robotics
Class 22	Brain-machine interface
Class 23	Human-machine interaction
Class 24	Creativity and AI
Class 25	Machine culture and AI as a cultural technology
Class 26	AI for scientific discovery

Tentative List of Course Topics and Materials

Class 1:

Introduction to Machine Intelligence

Topics / Questions

- Structure of the class, expectations of students
- Discuss:
 - How do you use AI?
 - What topics in AI interest you?
 - How does AI impact you?
 - Are you excited about AI, concerned about it, or both?
 - Overview of the field; what “intelligence” means in biological and artificial systems.

Readings and Materials

- Introduction of AI Snake Oil (2024) by Arvind Narayanan and Sayash Kapoor.
- Melanie Mitchell (2024). The metaphors of artificial intelligence. *Science* Vol 386, Issue 6723.
- Melanie Mitchell (2024). [“AI now beats humans at basic tasks.” Really?](#) Substack: AI: A Guide for Thinking Humans.
- One perspective on AI by experts:
 - AI will be smarter than humans in 20 years, 'Godfather of AI' warns. CUOMO. [Video](#)
 - Ilya: the AI scientist shaping the world. *The Guardian*. [Video](#)
- A different perspective:
 - Michael I. Jordan. Lecture at the "AI Action Summit" (Paris, 2025). [YouTube](#) (**Watch the first 9 minutes.**)
 - Emily Bender: How should regulators think about "AI"? [YouTube](#)

Class 2:

History of AI: Emergence from the fields of philosophy, mathematics, psychology, and neuroscience

Topics / Questions

- Tracing AI’s intellectual origins and interdisciplinary roots.
- What fields did AI emerge from?
- What were some of the main ideas?

Readings and Materials

- Eliza Strickland (2021). The turbulent past and uncertain future of AI. *IEEE Spectrum*.

- Daniel Crevier (1993). The first AI program: Defining the field. Ch. 2 in: The Tumultuous History of the Search for Artificial Intelligence. Basic Books.
- Daniel Crevier (1993). The dawn of the golden years: 1956-70. Ch. e in: The Tumultuous History of the Search for Artificial Intelligence. Basic Books.

Class 3:

Foundations and perspectives on machine intelligence

Topics / Questions

- Background and perspectives on AI
- Different approaches to AI.
- What is the purpose of AI?

Readings and Materials

- Michael I. Jordan (2019). Artificial intelligence – The revolution hasn't happened yet. *Harvard Data Science Review*.
- Yarden Katz (2012). [Noam Chomsky on where artificial intelligence went wrong](#). *The Atlantic*.

Class 4:

On intelligence, artificial general intelligence, and superintelligence

Topics / Questions

- Definitions and debates surrounding the scope and limits of machine intelligence.
- What is intelligence?
- What is artificial general intelligence (AGI)?
- What is superintelligence?

Readings and Materials

- Shane Legg, Marcus Hutter (2007). A collection of definitions of intelligence. *Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms*, 157, 17-24.
- Melanie Mitchell (2024). Debates on the nature of artificial general intelligence. *Science* Vol 383, issue 6689.
- Melissa Heikkilä (2025). Margaret Mitchell: artificial general intelligence is 'just vibes and snake oil'. *Financial Times*.
- Cade Metz (2025). Why We're Unlikely to Get Artificial General Intelligence Anytime Soon. *New York Times*.

Class 5:

Elements of an AI system: Data, algorithms, compute, and technology

Topics / Questions

- The architecture of modern AI and its enabling technologies.
- Discuss the main 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

Readings and Materials

- Introduction paragraph of AI Snake Oil (2024) by Arvind Narayanan and Sayash Kapoor.
- Why computer vision is a hard problem for AI (2024). *Quanta Magazine*. [YouTube](#)
- Anima Anandkumar (2018). [Trinity of artificial intelligence](#). TEDxIndianaUniversity. **Watch the first 11 minutes.**
- Data-centric AI Resource Hub. [Watch the two-minute video by Andrew Ng](#).
- Machine Learning Explained in 100 Seconds. [YouTube](#)
- Our World in Data (2023). [Charted](#)

Class 6:

Learning: 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, priors, or innate systems.

Readings and Materials

- Sutton R (2019). [The bitter lesson](#).
- Two Minute Papers. A bitter lesson - compute reigns supreme! [YouTube](#)
- Selections from Chapter 3 of *Genius Makers* and from interviews with Yoshua Bengio and Yann LeCun from *Architects of Intelligence* (2018), edited by Martin Ford. Packt Publishing.
- Rodney Brooks (2019). [A better lesson](#).
- Tenenbaum JB, Kemp C, Griffiths TL, Goodman ND (2011). How to grow a mind: statistics, structure, and abstraction. *Science* Vol 331: 1279-1285. **Read the first page until the end of the passage where the “three central questions” are presented.**
- Gary Marcus (2025). [The fever dream of imminent superintelligence Is finally breaking](#). *New York Times*, September 3, 2025.

Class 7:

Machine learning and artificial neural networks

Topics / Questions

- History of neural networks and their biological inspiration in brain science.

Readings and Materials

- Neural networks and deep Learning: Crash Course AI #3. [YouTube](#)
- Training Neural Networks: Crash Course AI #4. [YouTube](#)
- Samuel K. Moore, David Schneider, Eliza Strickland (2021). How deep learning works. *IEEE Spectrum*.
- [Neocognitron: A neural network model for a mechanism of visual pattern recognition](#) (1986).
- Cade Metz (2021). Promise. Chapter 2 in *Genius Makers: The Mavericks Who Brought AI to Google, Facebook, and the World*. Dutton.

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

Class 8:

Language and AI, part 1: How LLMs work

***** In-class Writing assignment 1**

Topics / Questions

- What are large language models (LLMs) and how do they work?
- Do LLMs "read" and "write"?
- Explore issues that arise when using generative AI tools like LLMs to assist in doing tasks, versus relying on such tools to create output which we then claim as our own.

Readings and Materials

- Before Siri and Alexa, there was ELIZA. [YouTube](#)
- Christopher Manning. Human language understanding and reasoning. *Daedalus* 151(2) Spring 2022.
- Visual Storytelling Team and Madhumita Murgia (2023). [Generative AI exists because of the transformer](#). *Financial Times*.
- Cho A, Kim GC, Karpekov A, Helbling A, et al. (2024). Transformer explainer: Interactive learning of text-generative models. *arXiv*. [Video](#). [Tool](#).
- John Pavlus (2025). When ChatGPT broke an entire field: An oral history. *Quanta Magazine*.

Class 9:

Language and AI, part 2: Meaning and reasoning in LLMs

Topics / Questions

- Discuss the relationship between LLMs and concepts like “meaning” and "reasoning."
- What LLMs capture (and miss) about semantics, reasoning, and understanding.

Readings and Materials

- Ellie Pavlick. Can large language models understand ‘meaning’? *Quanta Magazine* (2024). [YouTube](#)
- Blaise Agüera y Arcas (2022). Do large language models understand us? *Daedalus* 151(2) Spring 2022.
- CHM Live | The Great Chatbot Debate: Do LLMs Really Understand? [YouTube](#).
- Selections from the *Connectionists* email list, featuring discussion points by Paul Cisek and Geoffrey Hinton.

Class 10: Midterm 1

Class 11:

Emotion, Motivation, and Affective Computing

Topics / Questions

- How emotion and motivation shape cognition and behavior, and how AI attempts (or fails) to model affective processes.
- How affect influences cognition and how machines attempt to simulate emotional processes.

Readings and Materials

- Man, K., Damasio, A. Homeostasis and soft robotics in the design of feeling machines. *Nat Mach Intell* 1, 446–452 (2019).
- TBD

Class 12:

AI companions and mental health

Topics / Questions

- The integration of AI into mental health and wellness domains.
- Applications, risks, and psychological consequences of emotionally responsive AI.
- Should AI applications with implications for mental health be regulated?

Readings and Materials

- Emotional risks of AI companions demand attention. *Nat Mach Intell* 7, 981–982 (2025). <https://doi.org/10.1038/s42256-025-01093-9>
- De Freitas, J., Cohen, I.G. Unregulated emotional risks of AI wellness apps. *Nat Mach Intell* 7, 813–815 (2025). <https://doi.org/10.1038/s42256-025-01051-5>
- Shteynberg, G., Halpern, J., Sadovnik, A. *et al.* (2024). Does it matter if empathic AI has no empathy? *Nat Mach Intell* 6, 496–497.
- David Adam (2025). Supportive? Addictive? Abusive? How AI companions affect our mental health. *Nature*.
- Rachel Fieldhouse (2025). Can AI chatbots trigger psychosis? What the science says. *Nature*.
- Akiki, T.J., Williams, L.M., Wolfers, T. *et al.* Transforming psychiatry with computational and brain-based methods. *Nat Comput Sci* 5, 844–847 (2025).
- [A Teen Was Suicidal. ChatGPT Was the Friend He Confided In.](#) *New York Times*.
- [What My Daughter Told ChatGPT Before She Took Her Life.](#) *New York Times*.

Class 13:

Consciousness and AI

Topics / Questions

- Are current AI systems “slightly conscious”?
- What would it take for an AI system to be conscious or sentient?

Readings and Materials

- David Chalmers (2023). [Could a large language model be conscious?](#) *Boston Review*.
- [A Neuroscientist and a Philosopher Debate AI Consciousness: Michael Graziano and David Chalmers.](#)
- Blaise Agüera Y Arcas. Life, intelligence, and consciousness: A functional perspective. X Long Now.

Class 14:

Neuroscience and AI

Topics / Questions

- What are the connections between artificial intelligence and neuroscience?
- Reciprocal insights between brain research and artificial neural networks.
- What levels of neuroscience (cells, synapses, genes, circuits...), and what types of neurobiological information, representations, concepts, or constructs, are potentially useful for AI systems?

Readings and Materials

- [AI on the Brain: A Q&A with Richard Zemel](#)
- Zador AM (2019). [A critique of pure learning and what artificial neural networks can learn from animal brains](#). *Nature Communications*, Vol 10, article 3770.
- Zador A et al. (2023). [Catalyzing next-generation artificial intelligence through NeuroAI](#). *Nature Communications*, Vol 14, article 1597.
- Luo, X., Recharadt, A., Sun, G. *et al.* Large language models surpass human experts in predicting neuroscience results. *Nat Hum Behav* **9**, 305–315 (2025).

Class 15:

Computational Neuroscience and Representation

***** In-class Writing assignment 2**

Topics / Questions

- How the brain encodes, transforms, and uses information, and parallels to artificial neural representations.
- Deepens the link between neural data and model representations—ideal for students interested in both experimental and computational approaches.
- How neurons encode information; representational similarity between brains and models.

Readings and Materials

- David Marr's Tri-Level Hypothesis. In *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W. H. Freeman and Company.
- Population coding, predictive coding, distributed representations.

Class 16:

Neural Plasticity and Lifelong Learning

Topics / Questions

- How biological learning systems handle adaptation over time—an unsolved problem in machine intelligence.
- Biological plasticity, memory consolidation, and continual learning in humans vs. catastrophic forgetting in AI.

Readings and Materials

- Kudithipudi, D., Aguilar-Simon, M., Babb, J. *et al.* Biological underpinnings for lifelong learning machines. *Nat Mach Intell* **4**, 196–210 (2022).

Class 17:

Cognitive science and AI

Topics / Questions

- Human information processing, perception, reasoning, and parallels in machine systems.
- Nature vs. nurture, in psychology and AI.

Readings and Materials

- John Pavlus (2025). [‘World Models,’ an Old Idea in AI, Mount a Comeback](#). *Quanta Magazine*.
- Josh Tenenbaum (Nov 8, 2024). [Scaling intelligence the human way](#). IPAM at UCLA. Watch "the first half of the talk" up until about 23 minutes, 46 seconds.
- van Rooij, I., Guest, O., Adolfi, F. *et al.* Reclaiming AI as a theoretical tool for cognitive science. *Comput Brain Behav* 7, 616–636 (2024).
- Giosuè Baggio (2025). Could machine learning help to build a unified theory of cognition? *Nature*, July 29, 2025.

Class 18:

Developmental psychology and AI

Topics / Questions

- What can AI researchers learn from developmental psychology?
- Why are deep learning researchers so interested in child development?

Readings and Materials

- Elisabeth Spelke (2022). Core knowledge. In: *What Babies Know*. Oxford University Press.
- Karen Adolph Keynote, AI Symposium 2024. [YouTube](#).

Class 19:

Social Cognition and Theory of Mind in Humans and Machines

Topics / Questions

- How humans infer others’ mental states and current AI efforts to model social reasoning.

Readings and Materials

- Antonella Marchetti, Federico Manzi, Giuseppe Riva, Andrea Gaggioli, and Davide Massaro (2025). Artificial Intelligence and the Illusion of Understanding: A Systematic Review of Theory of Mind and Large Language Models. *Cyberpsychology, Behavior, and Social Networking*, Vol. 28, No. 7.

Class 20: Midterm 2

Class 21:

Embodied AI and robotics

Topics / Questions

- How does thinking about AI shift when the main purpose of the machine is to move in a 3D world, like a robot?
- How physical interaction grounds perception, learning, and intelligence.
- Areas within robotics, such as soft robotics, aerial and aquatic robotics, dexterous hand movements, artificial skin, locomotion, and feeling machines.
- Do we want feeling machines?

Readings and Materials

- Sami Haddadin (2025). [AI that moves, adapts, and learns: The future of embodied intelligence](#). Columbia AI
- Ken Goldberg (2024). Why don't we have better robots yet? TED. [YouTube](#)
- [MIT roboticists debate the future of robotics](#) (2025).
- “Data will solve robotics and automation: True or false?”: A debate. *Science Robotics* (2025).
- Rodney Brooks (2025). [Why Today’s Humanoids Won’t Learn Dexterity](#).

Class 22:

Brain-machine interface

Topics / Questions

- What is a brain-machine interface and how do they work?
- What are the main challenges of adding machine technology to a human brain and body?
- Direct neural interfaces, neuroprosthetics, and ethical considerations.

Readings and Materials

- New brain-computer interface (BCI) allows man with ALS to 'speak' again using brain implant and AI. [YouTube](#)
- Thought-controlled prosthetics: A brain-computer interface breakthrough (2024). [YouTube](#)
- Grégoire Courtine: The future of paralysis treatment (2023). [YouTube](#)
- Brain stimulation in treatment-resistant depression. [Text](#). [Videos](#)
- Gemma Conroy (2025). A mind-reading brain implant that comes with password protection. *Nature*.
- Miryam Naddaf (2024). Mind-reading devices are revealing the brain’s secrets. *Nature*.

Class 23:

Human-machine interaction

Topics / Questions

- How does human behavior and cognition change when interacting with machines?
- What happens when humans have their sense of selves disrupted by machine performance?

Readings and Materials

- Shoko Suzuki (2025). People are more likely to cheat when they delegate tasks to AI. *Nature*.
- AlphaGo documentary.

Class 24:

Creativity and AI

Topics / Questions

- Can machines be creative?
- In an era of AI, what does it mean to create something versus generate it?
- Is AI art good?

Readings and Materials

- **TBD**

Class 25:

Machine culture and AI as a cultural technology

***** In-class Writing assignment 3**

Topics / Questions

- Culture mediated or generated by machines; implications for psychology, science, and society.
- Future impact of machines on cultural evolution.
- AI as a new kind of cultural and social technology.

Readings and Materials

- Brinkmann, L., Baumann, F., Bonnefon, JF. *et al.* (2023). Machine culture. *Nat Hum Behav* **7**, 1855–1868.
- Henry Farrell, Alison Gopnik, Cosma Shalizi, James Evans. Large AI models are cultural and social technologies. *Science*, March 14, 2025.

Class 26:

AI for scientific discovery

Topics / Questions

- How is AI being used for scientific discovery?

Readings and Materials

- Xin, H., Kitchin, J.R. & Kulik, H.J. Towards agentic science for advancing scientific discovery. *Nat Mach Intell* **7**, 1373–1375 (2025). <https://doi.org/10.1038/s42256-025-01110-x>
- Messeri, L., Crockett, M.J. Artificial intelligence and illusions of understanding in scientific research. *Nature* **627**, 49–58 (2024). <https://doi.org/10.1038/s41586-024-07146-0>
- Nicola Jones (2025). AI ‘scientists’ joined these research teams: here’s what happened. *Nature*.
- Richard Van Noorden & Jeffrey M. Perkel (2023). AI and science: what 1,600 researchers think. *Nature*.
- Arvind Narayanan & Sayash Kapoor (2025). Why an overreliance on AI-driven modelling is bad for science. *Nature*.

FINAL EXAM

TBD

Class Policies

Enrollment. If the course is full, PhD students in the Psychology department, senior Psychology majors, senior Neuroscience and Behavior majors, and Psychology postbacs in the Certificate Program will have priority. Other things being equal, students who have the best preparation and strongest motivation will be selected.

Academic Integrity. Academic honesty includes presenting only your own work in exams and assignments, and correctly attributing others' ideas where Page 7 appropriate. Taking credit for work that is not your own is a serious violation within the academic community, and anyone found to be cheating or plagiarizing in this class will be reported to the university. Detailed definitions and examples of academic dishonesty (and a rundown of the consequences) are available in Columbia's Guide to Academic Integrity ([http://www.college.columbia.edu/academics integrity](http://www.college.columbia.edu/academics%20integrity)).

That said, I appreciate that the lines aren't always clear, so if you have any questions about how to properly cite a source or build upon others' ideas, or if you're feeling stressed out about the class workload (or about anything else), you should feel free to speak with me. Students with Disabilities. Students with special needs who may require accommodations should make an appointment to see me as soon as possible. If you have not already done so, stop by the Office of Disability Services (ODS) on the 7th floor of Lerner Hall to register for support services. ODS often requires two weeks to process an application, so please contact them as soon as you can, preferably before the course begins.