

**The AI Awakening:  
Implications for the Economy and Society**  
Econ295/CS323 | Spring 2025

**Subject to Change: Check back frequently | This Version: May 25, 2025**

**Instructor:** Erik Brynjolfsson

**TAs:** Marc-Aurele Stark ([mnstark@stanford.edu](mailto:mnstark@stanford.edu)) (Head TA); Alma Andino Frydman ([aandinof@stanford.edu](mailto:aandinof@stanford.edu)); Chiho Im ([chihoim@stanford.edu](mailto:chihoim@stanford.edu)); Matthew Merrigan ([merrigan@stanford.edu](mailto:merrigan@stanford.edu));

**Course Manager:** Susan Young

**Administrative Assistant:** Megan Deason ([mdeason@stanford.edu](mailto:mdeason@stanford.edu))

**Class Schedule and Locations:**

- Lectures: Tuesdays, 4:30-5:50pm, Gates Computer Science Building, Room B01
- Discussion: Thursdays, 4:30-5:50pm, Lathrop 018
- Final Presentations Sessions
  - Friday, June 6, 7-10pm ([our final exam slot](#)) in Mackenzie Huang 300

**Instructor Office Hours:** Tuesday 11am-12pm and by appointment at Gates 136 and via [RealAvatar](#)

**TA Office Hours:** Tuesday 1:00-2:00pm

**Canvas:** <https://canvas.stanford.edu/courses/209172>

**Introduction form:** <https://forms.gle/sqbZoHRt5YGFHx9M6>

## INTRODUCTION

The capabilities of AI have improved rapidly over the past few years, but our economic models, institutions, organizations, and skills have not kept up. In this growing gap lie many of society's greatest challenges and opportunities.

This course will explore how AI can and will transform our economy and society in the coming years. Each week, we will learn from frontier researchers and industry leaders in technology, economics, and business, read the relevant research, and discuss the implications. Small teams of students will work on AI-related projects and present them at the end of the term.

## CLASS ORGANIZATION

### Requirements

1. Class attendance at 12 sessions, including ten regular 80-minute meetings (Tuesdays, 4:30-5:50pm) plus one Project Workshop (Thursday April 24 or Thursday May 1) and one three-hour session (7-10pm on June 6) for final project presentations is required. Students are also strongly encouraged to attend all seven Thursday sessions.
2. Weekly readings: Typically 2-4 required readings and 4-8 optional readings
3. Assignments to be submitted on Canvas
  - a. Each week, several of the memo writers will discuss their memos in class.
  - b. Assignments are due by 9pm the Sunday before class, via [Canvas](#)
4. One final team project ([see guidelines](#))
  - a. Team formation by 5pm on April 12.
    - i. Teams of up to three people (business plans or policy proposals), and up to two people for research proposals.
    - ii. Teams should be diverse on multiple dimensions, with no more than two students from the same program.
  - b. Update 1: Progress update on final team project; due by 11.59pm on May 2
  - c. Update 2: First draft of final team project; due by 11.59pm on May 20
  - d. Team presentation with Q&A in class on June 6
  - e. Proposal for transformative AI venture (for profit or nonprofit), policy proposal, or research proposal; due by 11.59pm on June 3

### Grading

- 20% weekly assignments
- 20% mid-quarter progress updates on final team project
- 20% final presentation content and delivery
- 20% final project document
- 20% contribution to class discussions

## SCHEDULE

#	Lecture Date	Class Speaker (Tuesday 4:30-5:50pm; Gates B01)	Section Date	Discussion Section (Thursday 4:30-5:50; Lathrop 018)
1	April 1	<a href="#">Erik Brynjolfsson</a> The AI Awakening	April 3	Introductions
2	April 8	<a href="#">Daniela Rus</a>	April 10	Finding Your Project Group
3	April 15	<a href="#">Eric Horvitz</a>	April 17	A Career in AI Research
4	April 22	<a href="#">Eric Schmidt</a>	April 24	Final Project Workshop I
5	April 29	Erik Brynjolfsson	May 1	Final Project Workshop II
6	May 6	<a href="#">David Autor</a>	May 8	Larry Summers (Gates 119)
7	May 13	<a href="#">Ronnie Chatterji</a>	May 15	VC Pitch Workshop
8	May 20	<a href="#">Jack Clark</a>	May 22	Economic Impact of AI on Labor
9	May 27	<a href="#">Sendhil Mullainathan</a>	May 29	TA Feedback Session
10	June 3	<a href="#">Mira Murati</a>		
11	June 6	<b>7-10pm</b> Final presentations: Research & Policy Proposals + Business Plans In Jen-Hsun Huang Engineering Center		

## Thursday Discussion Section Calendar

### **Week 1: Introductions**

Thursday April 3

This is your first opportunity to introduce yourself to your peers in the class and share your goals for the quarter. You will share a brief bio and topics you would like to explore through your final project. Connect with your classmates and begin meeting possible teammates!

### **Week 2: Finding Your Final Project Group**

Thursday April 10

This structured workshop will help you get to know your classmates deeper and form final project groups around shared interests. We will hold “speed dating” rounds to allow you to have individual conversations, followed by small group discussions sorted by topic and project type.

### **Week 3: A Career in AI Research**

Thursday April 17

What are the most pressing questions in the field of AI research? What problems are researchers solving at leading AI companies and how did they get to their role? Come hear top researchers present their work across model foundations, alignment and safety, and societal impacts. Speakers will give short presentations on their groundbreaking work and will then host small roundtable discussions.

### **Week 4: Feedback on Final Project, Section I**

Thursday April 24

Each project team will be matched with two other groups to provide feedback during this session. Team members should complete the feedback form before their assigned Thursday workshop, so they can use the section for a constructive discussion. Paired groups will have 40 minute sessions—20 minutes per team—to exchange feedback, practice presenting, and get TA input. Teams will only need to attend ONE of the two workshops.

### **Week 5: Feedback on Final Project, Section II**

Thursday May 1

Each project team will be matched with two other groups to provide feedback during this session. Team members should complete the feedback form before their assigned Thursday workshop, so they can use the section for a constructive discussion. Paired groups will have 40 minute sessions—20 minutes per team—to exchange feedback, practice presenting, and get TA input. Teams will only need to attend ONE of the two workshops.

### **Week 6: Larry Summers (Gates 119)**

Thursday May 8

Lawrence H. Summers is a leading American economist and public official known for his influential roles in academia and government. He served as Chief Economist at the World Bank from 1991 to 1993, U.S. Treasury Secretary under President Clinton from 1999 to 2001, and Director of the National Economic Council for President Obama from 2009 to 2010. From 2001 to 2006, he was President of Harvard University. A Harvard-trained economist, Summers has long shaped global economic discourse through his work on fiscal policy, globalization, and

financial crises. Today, he remains an active voice in economic commentary and research as the Charles W. Eliot University Professor and President Emeritus at Harvard University.

**Week 7: How do you pitch an AI startup to VCs?** Thursday May 15

VC firms invested about \$130 billion into AI start-ups in 2024. What kinds of AI ventures secure funding, and which ones fail to impress? Our pitch workshop will bring together leading VCs and give students a chance to receive guidance on their ideas and presentation. Groups should come prepared with a 2-minute pitch and a minimum of two team members in attendance.

**Week 8: Economic Impact of AI on Labor** Thursday May 22

This workshop will explore the potential for rapid economic growth enabled by advanced AI in production and in research. Phil Trammel, a postdoc at the Digital Economy Lab, will dive into the nitty-gritty of three scenarios: (i) full automation of production only, (ii) full automation of research only, and (iii) partial automation of both. The workshop will focus on mathematical formulations of AI in macroeconomic growth models.

**Week 9: TA Feedback Session** Thursday May 29

Looking for more feedback on your project? Drop by this TA feedback workshop to cowork on your final with other teams and get support from the teaching team in preparation for your final presentation.

## **Tuesday Classes**

### **April 1: Erik Brynjolfsson**

Main readings:

1. Sutton, R. (2019) "[The Bitter Lesson](#)" *Incomplete Ideas*, March 13.
2. Brynjolfsson, E. & Unger, G. (2023,) "[The macroeconomics of artificial intelligence.](#)" *IMF Finance and Development Magazine*.
3. Brynjolfsson, E., & Mitchell, T. (2017). [What can machine learning do? Workforce implications.](#) *Science*, 358(6370), 1530-1534.

Optional readings:

4. Beaver, C. (2023). [ChatGPT broke the Turing test — the race is on for new ways to assess AI.](#) *Nature*, 619(686-689).
5. Turing, A. (1950). [Computing Machinery and Intelligence.](#)
6. Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., ... & Amodei, D. (2020). [Scaling laws for neural language models.](#) *arXiv preprint arXiv:2001.08361*.
7. Brynjolfsson, E., Thierer, A., & Acemoglu, D. (2024, March). [Navigating the Future of Work: Perspectives on Automation, AI and Economic Prosperity.](#) *AEI*.
8. Brynjolfsson, E., Li, D., & Raymond, L. (2023, April). [Generative AI at Work.](#) *National Bureau of Economic Research Working Paper 31161*.
9. Brynjolfsson, E., Rock, D., & Syverson, C. (2018). [Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics.](#) In *The economics of artificial intelligence: An agenda* (pp. 23-57). University of Chicago Press.
10. McCarthy, J, M. L. Minsky, N. Rochester, C.E. Shannon (1955) "[A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence](#)".

Assignment:

- Complete [the introduction form](#) to help with team formation. Share your background and interests. Tell us your 3 main goals for the course.
- Complete the [speaker dinner preferences form](#). We've asked all the speakers to stay for dinner, but they have not all confirmed and their plans may change.

### **April 8: Daniela Rus**

Main readings:

1. Rus, D. and G. Mone, (2024) [The robots are coming. And that's a good thing.](#) *Technology Review*.
2. Rus, D. (2024). [How AI will step off the screen and into the real world](#) [Video]. *TED. YouTube*.
3. Knight, W. (2024). [Liquid AI is redesigning the neural network.](#) *WIRED*.

4. [Physical Intelligence Blog](#) (some great videos of SOTA robotics)
5. Rus, D. (2025). [To interact with the real world, AI will gain physical intelligence.](#) *WIRED*.

Optional readings:

1. Perrault, R., Y. Gil, et al. [The 2025 AI Index Report](#), Stanford HAI.
2. Nietzel, M. (2021). [Meet the world's most influential women engineers.](#) *Forbes*.
3. Wiggers, K. (2023). [Liquid AI, a new MIT spinoff, wants to build an entirely new type of AI.](#) *TechCrunch*.
4. Semuels, A. (2024). [How automation is helping companies prepare for labor shortages.](#) *Time*.
5. Rus, D. (2024). [To compete with China on AI, we need a lot more power.](#) *Washington Post*.

### April 15: Eric Horvitz

Main readings:

1. E. Horvitz and T. Mitchell, [Scientific Progress in Artificial Intelligence: History, Status, and Futures](#). In: Realizing the Promise and Minimizing the Perils of Artificial Intelligence for the Scientific Community. Kathleen Hall Jamieson, Anne-Marie Mazza, and William Kearney (Eds.) (2024). University of Pennsylvania Press.
2. Goh E, Gallo R, Hom J, et al. [Large Language Model Influence on Diagnostic Reasoning: A Randomized Clinical Trial.](#) *JAMA Netw Open*. 2024;7(10):e2440969. doi:10.1001/jamanetworkopen.2024.40969
3. Brynjolfsson, E. (2022). [The Turing Trap: The promise & peril of human-like artificial intelligence.](#) *Daedalus*, 151(2), 272-287.

Optional readings:

1. H. Nori, N. Usuyama, N. King, S. Mayer McKinney, X. Fernandes, S. Zhang, **E. Horvitz**. [From Medprompt to o1: Exploration of Run-Time Strategies for Medical Challenge Problems and Beyond](#), November 2024.
2. S. Bubeck, V. Chandrasekaran, R. Eldan, J. Gehrke, E. Horvitz, E. Kamar, P. Lee, Y. Lee, Y. Li, S. Lundberg, H. Nori, H. Palangi, M. Ribeiro, Y. Zhang, [Sparks of Artificial General Intelligence: Early experiments with GPT-4](#), arXiv preprint arXiv:2303.12712
3. Melanie Mitchell, [The Turing Test and our shifting conceptions of intelligence.](#) *Science* **385**,eadq9356(2024).DOI:10.1126/science.adq9356
4. Bieber, C. (2023). [ChatGPT broke the Turing test — the race is on for new ways to assess AI.](#) *Nature*, 619(686-689).

### April 22: Eric Schmidt

Main readings:

1. Hendrycks, D., Schmidt, E., & Wang, A. (n.d.). [Superintelligence strategy](https://nationalsecurity.ai). Retrieved March 11, 2025, from <https://nationalsecurity.ai>
2. Schmidt, E. (Host). (2023). Eric Schmidt on the risks and rewards of artificial intelligence (No. 3) [Audio podcast episode]. In Possible. Spotify. <https://open.spotify.com/episode/5NQFPbINw8ewxKollDpiYN>
3. Schmidt, E.; Raghu, M.; & Zaharia, M. (2023) – [“Does One Large Model Rule Them All? Predictions on the Future AI Ecosystem.”](#)

Optional readings:

4. Schmidt, E. (2022) – [“AI, Great Power Competition & National Security.”](#)
5. Schmidt, E. (2023) – [“Innovation Power: Why Technology Will Define the Future of Geopolitics.”](#) – Foreign Affairs.
6. Kissinger, H., Schmidt, E., & Huttenlocher, D. (2023, February 24). [ChatGPT Heralds an Intellectual Revolution](#)

**April 29: Erik Brynjolfsson (for Mira Murati)**

Main readings:

1. [“OpenAI’s Former Chief Technology Officer Starts Her Own Company”](#), NYT
2. Things Mira has led at OpenAI
  - a. [GPT-4 Technical Report](#)
  - b. [What is o1? What is a reasoning model?](#) How are they different from GPT-4?
3. “Thinking” Machines
  - a. Manning, C. D. (2022). [Human language understanding & reasoning](#). *Daedalus*, 151(2), 127-138.
  - b. Legg, S., & Hutter, M. (2007). Universal intelligence: A definition of machine intelligence. *Minds and machines*, 17, 391-444.
  - c. [On the Measure of Intelligence](#) - <https://arxiv.org/abs/1911.01547>

Optional readings:

4. Video generation models as world simulators, [OpenAI \(2024\)](#)
5. Our approach to AI safety [\(OpenAI 2023\)](#)
6. Moving AI governance forward [\(OpenAI 2023\)](#)
7. How OpenAI is approaching worldwide elections [\(OpenAI 2024\)](#)
8. Tech Accord to Combat Deceptive Use of AI in 2024 Elections [\(AI Elections Accord 2024\)](#)

**May 6: David Autor**

#### Main readings:

1. National Academies of Sciences, Engineering, and Medicine. (2024). Artificial intelligence and the future of work. The National Academies Press.  
<https://doi.org/10.17226/27644>
2. Autor, D. (2024) – “[AI Could Actually Help Rebuild the Middle Class.](#)” – Noema Magazine.
3. Autor, D. (2015) – “[Why Are There Still So Many Jobs? The History and Future of Workplace Automation.](#)” – J. of Economic Perspectives, 29(3): 3–30.
4. Narayanan, Arvind & Kapoor, Sayash (2025) “[AI as Normal Technology.](#)” April.
5. Schwarcz, D., Manning, S., Barry, P., Cleveland, D. R., Prescott, J. J., & Rich, B. (2025). [AI-powered lawyering: Ai reasoning models, retrieval augmented generation, and the future of legal practice.](#)”

#### Optional readings:

6. Eloundou, T., et al. (2023) – “[GPTs are GPTs: An Early Look at the Labor Market Impact of Large Language Models.](#)” – OpenAI/OpenResearch working paper.
7. Aghion, P., Jones, B. F., & Jones, C. I. (n.d.). Artificial intelligence and economic growth. National Bureau of Economic Research. Retrieved March 11, 2025, from <https://www.nber.org/system/files/chapters/c14015/c14015.pdf>
8. Acemoglu, D. & Restrepo, P. (2018) – “[Artificial Intelligence, Automation, and Work.](#)” In The Economics of AI: An Agenda (NBER chapter).
9. Humlum, Anders, and Vestergaard, Emilie. (2025), “[Large Language Models, Small Labor Market Effects](#)”

#### May 13: Ronnie Chatterji

#### Main readings:

1. [Dr. Ronnie Chatterji named OpenAI's first Chief Economist](#)
2. [OpenAI's Economic Blueprint](#)
3. “How to Allocate \$52 Billion in CHIPS Funding”
  - a. <https://pitchforeconomics.com/episode/the-chips-act-explained-with-ronnie-chatterji/>
  - b. <https://www.statecraft.pub/p/how-to-allocate-52-billion-in-chips>

#### Optional readings:

4. [Where has the CHIPS act money gone.](#) (The Verge, 2024)
5. [Governing With Artificial Intelligence: Are Governments Ready?](#) (OECD, 2024)
6. OECD (2023), “[A blueprint for building national compute capacity for artificial intelligence.](#)”, *OECD Digital Economy Papers*, No. 350, OECD Publishing, Paris, <https://doi.org/10.1787/876367e3-en>.

7. [AI's Power Requirements Under Exponential Growth](#), (RAND, 2025)
8. [To compete with China on AI, we need a lot more power](#) (Washington Post, 2024)
9. [Biden's microchip man says goodbye](#), (POLITICO, 2023)
10. [The Role of Innovation and Entrepreneurship in Economic Growth](#)

## May 20: Jack Clark

### Main readings:

1. [Constitutional AI: Harmlessness from AI Feedback \(Anthropic, 2022\)](#)  
Introduces Anthropic's core method for training helpful and harmless models using written principles instead of human demonstrations—a cornerstone of Claude's alignment.
2. [Tracing the Thoughts of a Large Language Model \(Anthropic, 2025\)](#)  
Explores Claude's internal reasoning via interpretability techniques, showing how it plans ahead and recognizes unsafe requests—perfect for showcasing frontier safety research.
3. [Import AI 410 \(Clark\)](#) - "Eschatological AI Policy Is Very Difficult"
4. [Import AI 405 \(Clark\)](#) - "What if we're right about AI timelines? What if we're wrong?"
5. [Anthropic's Economic Index: AI's Impact on Software Development](#)
6. [Anthropic's OSTP RFI Filing](#)

### Optional readings:

1. [Collective Constitutional AI \(Anthropic, 2023\)](#) Describes an experiment aligning a model with principles sourced from public input—raises great questions around whose values AI should reflect.
2. [Mapping the Mind of a LLM \(Anthropic, 2024\)](#) Shows how Claude 3 encodes human-interpretable concepts in its neurons—foundational work in interpretability.
3. [Emergent Abilities of LLMs \(Wei et al., 2022\)](#) Shows how unexpected capabilities emerge as models scale—critical for understanding why Anthropic is cautious with scaling.
4. [The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation \(Clark et al., 2018\)](#) One of the earliest major warnings about how AI could be misused in cyberattacks, misinformation, and surveillance

## May 27: Sendhil Mullainathan

### Main readings:

1. Mullainathan, S. (2020). [Sendhil Mullainathan on how AI can help counter human bias](#). [Video]. *Chicago Booth Review*. Youtube.
2. Mullainathan, S., & Speiss, J. (2017). [Machine Learning: An Applied Econometric Approach](#). *Journal of Economic Perspectives* 31 (2): 87–106.
3. Cockrell, J. (2023). [What's the next big question in the social sciences? Ask an algorithm](#). *Chicago Booth Review*.

Optional readings:

1. [Nightingale Open Science](#).
2. Pierson, E., et al. (2025). [Using Large Language Models to Promote Health Equity](#). *NEJM AI* 2 (2): A1p2400889.

### June 3: Mira Murati

Main readings (review from April 29)

1. [“OpenAI’s Former Chief Technology Officer Starts Her Own Company”](#), NYT
2. Things Mira has led at OpenAI
  - a. [GPT-4 Technical Report](#)
  - b. [What is o1? What is a reasoning model?](#) How are they different from GPT-4?
3. “Thinking” Machines
  - a. Manning, C. D. (2022). [Human language understanding & reasoning](#). *Daedalus*, 151(2), 127-138.
  - b. Legg, S., & Hutter, M. (2007). Universal intelligence: A definition of machine intelligence. *Minds and machines*, 17, 391-444.
  - c. [On the Measure of Intelligence](#) - <https://arxiv.org/abs/1911.01547>

Optional readings:

4. Video generation models as world simulators, [OpenAI \(2024\)](#)
5. Our approach to AI safety [\(OpenAI 2023\)](#)
6. Moving AI governance forward [\(OpenAI 2023\)](#)
7. How OpenAI is approaching worldwide elections [\(OpenAI 2024\)](#)
8. Tech Accord to Combat Deceptive Use of AI in 2024 Elections [\(AI Elections Accord 2024\)](#)

### June 6: Final Presentations (7-10pm in Jen-Hsun Huang Engineering Center)