

# The AI Awakening: Implications for the Economy and Society

Econ 295/CS323 | Spring 2023

*Subject to Change: Check back frequently*

Instructor: Erik Brynjolfsson

Co-instructors: Sebastian Thrun and James Manyika

TA: Ruyu Chen ([ruyuchen@stanford.edu](mailto:ruyuchen@stanford.edu)), Rebecca Janßen ([rjanssen@stanford.edu](mailto:rjanssen@stanford.edu))

Course Manager: Susan Young

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

**Class Meeting Time:** Tuesdays from 4:30-5:50pm, April 4 - June 6

**Location:** Gates Building, Room B03 (note new room)

**Optional Unsupervised Learning Sessions:** Four selected Thursdays from 4:30-5:30pm at Gates B03 (April 6, April 20, and the last two sessions TBC)

**Instructor Office Hours:** Monday 11:00am-12:00noon at Gates 136

**TA Office Hours:** Tuesday 1:00-2:00pm at Gates 138

**Canvas:** <https://canvas.stanford.edu/courses/174140>

## INTRODUCTION

Our economic institutions, organizations, and skills have not kept up with rapid advances in AI. In this growing gap lie many of society's greatest challenges and opportunities.

This course will explore how the advances in 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 and economics, read the relevant research, and discuss the implications.

## CLASS ORGANIZATION

### Requirements

1. Class attendance at 10 sessions, including nine regular 80-minute meetings plus one four-hour session (4:30-8:30pm on June 6) for final project presentations.
2. Weekly readings: Typically 2-3 required readings and 3-5 optional readings
3. Six (of eight) one-page memos due by 5pm the *Sunday before class* via Canvas
  - a. Each week, several of the memo writers will be asked to discuss their memos in class and some may be invited to join a private dinner with that week's guest speaker.
4. One final team project ([see guidelines](#))
  - a. Team formation by 5pm on Friday, April 21. Teams of up to four people should be diverse on multiple dimensions, with no more than two students from the same program.
  - b. Progress update on final team project; due by 5pm on Friday, May 12
  - c. Team presentation with Q&A in class on June 6
  - d. 5-page proposal for transformative AI venture (for profit or nonprofit) or policy proposal; due by 5pm on June 9 (details to come)

## Grading

- 24% six (of eight) one-page memos or assignments
- 10% mid-quarter progress update on final team project
- 15% final presentation content and delivery
- 25% final project document
- 26% contribution to class discussions

## SCHEDULE

#	Date	Time	Topic	Guests (subject to change)
1	April 4	4:30-5:50pm	Introduction: The AI Awakening	<a href="#">Erik Brynjolfsson</a> <a href="#">Sebastian Thrun</a> <a href="#">Jack Clark</a>
2	April 11	4:30-5:50pm	Large language models	<a href="#">Mira Murati</a>
3	April 18	4:30-5:50pm	Technical progress in AI	<a href="#">Jeff Dean</a>
4	April 25	4:30-5:50pm	Opportunities and Risks of AI	<a href="#">Eric Schmidt</a>
5	May 2	4:30-5:50pm	Human + machine	<a href="#">Alexandr Wang</a>
6	May 9	4:30-5:50pm	Business Implications of AI	<a href="#">Bindu Reddy</a>
7	May 16	4:30-5:50pm	Work and employment	<a href="#">Laura Tyson</a>
8	May 23	4:30-5:50pm	A World without Work	<a href="#">Mustafa Suleyman</a>
9	May 30	4:30-5:50pm	AI, Geopolitics and Democracy	<a href="#">Condoleezza Rice</a>
10	June 6	4:30-8pm	Final projects presentations	<a href="#">Roy Bahat</a> <a href="#">Songye Yoon</a>

## READINGS

See [Daedalus special issue on AI and Society](#) (Volume 151 Issue 2) for many of the readings

### April 4: Intro and overview; Progress in AI

Main readings:

1. Manyika, James. "[Getting AI right: Introductory notes on AI & society.](#)" *Daedalus* 151.2 (2022): 5-27.
2. Clark, Jack. [Import AI 321: Open source GPT-3; giving away democracy to AGI companies; GPT-4 is a political artifact.](#) March 21, 2023
3. Wolfram, Stephen. "[What is ChatGPT doing... And Why Does it Work?](#)" *Writings*, February 14, 2023. (skim)

Optional readings:

4. Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. "[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, (2018).
5. Ganguli, D., Hernandez, D., Lovitt, L., Askell, A., Bai, Y., Chen, A., ... & Clark, J. (2022, June). "[Predictability and surprise in large generative models](#)". In *2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1747-1764).
6. Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson. "[Firm investments in artificial intelligence technologies and changes in workforce composition.](#)" Available at SSRN (2022).
7. Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, "[Prediction Machines: The Simple Economics of Artificial Intelligence](#)", Harvard Business Press, 2018.
8. Shadbolt, Nigel. "[From So Simple a Beginning": Species of Artificial Intelligence.](#)" *Daedalus* 151, no. 2 (2022): 28-42.
9. Zhang, Daniel, Saurabh Mishra, Erik Brynjolfsson, John Etchemendy, et al. "[The AI Index 2022 Annual Report.](#)" AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University, March 2022.

### April 11: Large language models/Generative models

Main readings:

1. Murati, Ermira. "[Language & coding creativity.](#)" *Daedalus* 151, no. 2 (2022): 156-167.
2. Bubeck, Sébastien, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee et al. "[Sparks of Artificial General Intelligence: Early experiments with GPT-4.](#)" *arXiv preprint arXiv:2303.12712* (2023).

Optional readings:

3. Manning, Christopher D. "[Human language understanding & reasoning.](#)" *Daedalus* (2022) 151, no. 2: 127-138.
4. Srivastava, Aarohi, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown et al. "[Beyond the imitation game: Quantifying and extrapolating the capabilities of language models.](#)" *arXiv preprint arXiv:2206.04615* (2022).
5. Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. "[Bert: Pre-training of deep bidirectional transformers for language understanding.](#)" *arXiv preprint arXiv:1810.04805* (2018).
6. Radford, Alec, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. "[Improving language understanding by generative pre-training.](#)" (2018).
7. Chen, Mark, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. "[Generative pretraining from pixels.](#)" In *International conference on machine learning*, pp. 1691-1703. PMLR, 2020.

#### **April 18: Technical progress in AI**

Main reading:

1. Dean, Jeffrey. "[A golden decade of deep learning: Computing systems & applications.](#)" *Daedalus* 151, no. 2 (2022): 58-74.
2. Kevin Scott. "[I Do Not Think It Means What You Think It Means: Artificial Intelligence, Cognitive Work & Scale](#)". *Daedalus* 2022; 151 (2): 75–84.

Optional readings:

3. Taddy, Matt. "[The technological elements of artificial intelligence.](#)" In *The economics of artificial intelligence: An agenda*, pp. 61-87. University of Chicago Press, 2018.
4. Stuart Russell and Peter Norvig. [Artificial Intelligence: A Modern Approach](#), 4th Edition
5. Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. "[Attention is all you need.](#)" *Advances in neural information processing systems* 30 (2017).
6. LeCun, Y., Bengio, Y. & Hinton, G. "[Deep learning](#)". *Nature* 521, 436–444 (2015).
7. OpenAI. "[GPT-4 Technical Report](#)". March 2023
8. Taddy, Matt. "[The technological elements of artificial intelligence.](#)" In *The economics of artificial intelligence: An agenda*, pp. 61-87. University of Chicago Press, 2018.
9. Wei, Jason, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zopf, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Lian, Jeff Dean, and William Fedus. "[Emergent Abilities of Large Language Models](#)". *Transactions on Machine Learning Research* (08/2022).
10. Narang, Sharan, and Aakanksha Chowdhery, "[Pathways Language Model \(PaLM\): Scaling to 540 Billion Parameters for Breakthrough Performance](#)". Google Research Blog.

## April 25: Opportunities and Risks of AI

Main reading:

1. Schmidt, Eric. "[AI, Great Power Competition & National Security.](#)" *Daedalus* 151, no. 2 (2022): 288-298.
2. Henry Kissinger, Eric Schmidt and Daniel Huttenlocher "[ChatGPT Heralds an Intellectual Revolution](#)" *WSJ*. Feb 24, 2023
3. Schmidt, Eric "[The Mystery Ahead of Us for Mankind](#)", Special Competitive Studies Project, April 6, 2023.
4. Klein, Ezra "[The Surprising Thing AI Engineers Will Tell You if You Let Them](#)". *The New York Times*, April 16, 2023.

Optional readings:

5. Bommasani, Rishi, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein et al. "[On the opportunities and risks of foundation models.](#)" *arXiv preprint arXiv:2108.07258* (2021). [Skim]
6. Kissinger, Henry A., Schmidt, Eric, & Huttenlocher, Daniel (2021). *The age of AI: and our human future*. Hachette UK.
7. Raghu, M., & Schmidt, E. (2020). A survey of deep learning for scientific discovery. *arXiv preprint arXiv:2003.11755*.

## May 2: Human + Machine; Reinforcement Learning from Human Feedback (RLHF)

Main reading:

1. Brynjolfsson, Erik. "[The Turing Trap: The promise & peril of human-like artificial intelligence.](#)" *Daedalus* 151, no. 2 (2022): 272-287.
2. Lipchik, Saul, "[The Next Elon Musk? Meet Alexandr Wang, CEO of Scale AI](#)" *South China Morning Post*, May 31, 2022.
3. Dickson, Ben "[What is reinforcement learning from human feedback \(RLHF\)?](#)" *TechTalks*, January 16, 2023.

Optional readings:

4. Ouyang, Long, Jeff Wu, Xu Jiang, et al., "[Training Language Models to Follow Instructions with Human Feedback](#)," OpenAI, 2022
5. Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb. "[Artificial intelligence: the ambiguous labor market impact of automating prediction.](#)" *Journal of Economic Perspectives* 33, no. 2 (2019): 31-50.
6. Korinek, Anton, and Joseph E. Stiglitz. "[Artificial intelligence and its implications for income distribution and unemployment.](#)" In *The economics of artificial intelligence: An agenda*, pp. 349-390. University of Chicago Press, 2018.

7. Acemoglu, Daron, and Pascual Restrepo. "[Artificial intelligence, automation, and work.](#)" In *The economics of artificial intelligence: An agenda*, pp. 197-236. University of Chicago Press, 2018.

## May 9: Business Implications of AI

Main reading:

1. Edd Gent (2023). The AI Evolution: How AI is Now Building Itself. Retrieved from <https://www.science.org/content/article/artificial-intelligence-evolving-all-itself>
2. The New York Times (March 2023). What's the Future for AI? Retrieved from <https://www.nytimes.com/2023/03/31/technology/ai-chatbots-benefits-dangers.html>
3. Quanta Magazine (January 2022). Researchers Build AI that Builds AI. Retrieved from <https://www.quantamagazine.org/researchers-build-ai-that-builds-ai-20220125/>

Optional readings:

4. Knight, Will, "[ChatGPT's most charming trick is also its biggest flaw](#)" *Wired*, December 7, 2022
5. Chui, Michael, James Manyika, Mehdi Miremadi, Nicolaus Henke, Rita Chung, Pieter Nel, and Sankalp Malhotra. "[Notes from the AI frontier: Insights from hundreds of use cases.](#)" *McKinsey Global Institute* 2 (2018).
6. Brynjolfsson, E., & McAfee, Andrew. (2017). [The Business of Artificial intelligence.](#) *Harvard Business Review*, 1, 1-31.
7. Varian, Hal. "[Artificial intelligence, economics, and industrial organization.](#)" In *The economics of artificial intelligence: an agenda*, pp. 399-419. University of Chicago Press, 2018.
8. Daniela Rus. "[The Machines from Our Future](#)". *Daedalus* 2022; 151 (2): 100–113.

## May 16: Work and employment

Main readings:

1. Laura D. Tyson, John Zysman; "[Automation, AI & Work](#)". *Daedalus* 2022; 151 (2): 256–271.
2. Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. "[GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models](#)". arXiv preprint arXiv:2303.10130. (2023)

Optional readings:

3. Spence, Michael. "[Augmentation, Value Creation & the Distribution of Income & Wealth](#)". *Daedalus* 2022; 151 (2): 244–255.
4. Brynjolfsson, Erik, and Tom Mitchell. "[What can machine learning do? Workforce implications.](#)" *Science* 358, no. 6370 (2017): 1530-1534.

5. Petroni, Fabio, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H. Miller, and Sebastian Riedel. "[Language models as knowledge bases?](#)." *arXiv preprint arXiv:1909.01066* (2019).

### **May 23: A World without Work**

Main reading:

1. Mustafa Suleyman. "[Four Features of the Coming Wave](#)". Chapter 7 of *The Coming Wave*. Please note: This document is confidential and must not be shared.
2. Altman, Sam. "[Moore's law for everything](#)"
3. Korinek, Anton, and Megan Juelfs. "[Preparing for the \(non-existent?\) Future of Work](#)". [No. w30172](#). National Bureau of Economic Research, 2022.

Optional readings:

4. Nordhaus, William D. "[Are we approaching an economic singularity? Information technology and the future of economic growth](#)." *American Economic Journal: Macroeconomics* 13, no. 1 (2021): 299-332.
5. Korinek, Anton, and Joseph E. Stiglitz. "[Artificial intelligence, globalization, and strategies for economic development](#)." No. w28453. National Bureau of Economic Research, 2021.
6. Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones. "[Artificial intelligence and economic growth](#)." In *The economics of artificial intelligence: An agenda*, pp. 237-282. University of Chicago Press, 2018.
7. Korinek, Anton. "[Language Models and Cognitive Automation for Economic Research](#)". No. w30957. National Bureau of Economic Research, 2023.

### **May 30: AI, Geopolitics and Democracy**

Main reading:

1. Carter, Ash. "[The Moral Dimension of AI-Assisted Decision-Making: Some Practical Perspectives from the Front Lines](#)". *Daedalus* 2022; 151 (2): 299–308.
2. Schmidt, Eric. "Innovation Power Why Technology Will Define the Future of Geopolitics" *Foreign Affairs*

Optional readings:

3. Beraja, Martin, Andrew Kao, David Y. Yang, and Noam Yuchtman. "[AI-tocracy](#)". forthcoming, *Quarterly Journal of Economics*
4. Brianna Nicker. "[The geopolitics of AI and the rise of digital sovereignty](#)". *Brookings News*, December 8, 2022