

**The AI Awakening:
Implications for the Economy and Society**

Econ 295/CS323 | Spring 2024

Subject to Change: Check back frequently | Last update: May 15, 2024

Instructor: Erik Brynjolfsson

TAs: James Da Costa (jdacosta@stanford.edu), Houda Nait El Barj (hnait@stanford.edu),
Marshall Mo (erxiaomo@stanford.edu)

Course Manager: Susan Young

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

Class Schedule and Locations:

Tuesdays 4:30-5:50pm, April 2-June 4 in Gates Computer Science Building, Room B03

Friday, June 7 7-10pm ([our final exam slot](#)) in Mackenzie Huang 300

Optional Sessions: Tuesdays 6-7pm unless otherwise noted, April 2-June 4

Location: Varies - Gates Computer Science Building, Room 119 and ChEM-H E153 (see schedule below for specifics)

Instructor Office Hours: Tuesday 11am-12pm and by appointment at Gates 136

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

Canvas: <https://canvas.stanford.edu/courses/191381>

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 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.

CLASS ORGANIZATION

Requirements

1. Class attendance at 11 sessions, including ten regular 80-minute meetings (Tuesdays, 4:30-5:50pm) plus one three-hour session (7-10pm on June 7) for final project presentations.
2. Weekly readings: Typically 2-4 required readings and 4-6 optional readings
3. One-page memos to be submitted on Canvas
 - a. Each week, several of the memo writers will be asked to discuss their memos in class.
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 up to two students from the same program.
 - b. Update 1: Progress update on final team project; due by 4pm on April 30
 - c. Update 2: First draft of final team project; due by 4pm on May 21
 - d. Team presentation with Q&A in class on June 4 or 7 (depending on which project type you select)
 - e. Proposal for transformative AI venture (for profit or nonprofit), policy proposal, or research proposal; due by 5pm on June 5

Grading

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

SCHEDULE

#	Date	Class Speaker + Topic (4:30-5:50pm)	Optional Session + Topic (6-7pm)
1	April 2	Erik Brynjolfsson The AI Awakening	Erik Brynjolfsson Student Introductions Gates 119
2	April 9	Eric Schmidt The Age of AI	Houda, James, and Marshall Team formation and project advice Gates 119
3	April 16	Mira Murati Large Language Models and Beyond	Houda, James, and Marshall Deeper dive into AI ChEM-H E153
4	April 23	Reid Hoffman Business of AI	Panos Madamopoulos-Moraris with Sarah Guo, Gonzague de Pirey, Omar Shaya Understanding the AI Landscape Gates 119
5	April 30	Anima Anandkumar AI and Creativity	Russell Wald Policy ChEM-H E153
6	May 7	James Manyika AI and Society	Houda, James, and Marshall Deeper dive into AI Gates 119
7	May 14	Shivon Zilis Augmenting Humans	Sandy Pentland Tech-based entrepreneurship ChEM-H E153
8	May 21	Lawrence Lessig AI and Democracy	Houda, James, and Marshall Project check-ins ChEM-H E153
9	May 28	Daniel Susskind A World Without Work	No optional session
10	June 4	Final presentations: Research and Policy Proposals	No optional session
11	June 7	*7-10pm* Final presentations: Business Plans Mackenzie Huang 300	

April 2: Intro and overview; Progress in AI and some implications

Main readings:

1. Brynjolfsson, E., & Mitchell, T. (2017). [What can machine learning do? Workforce implications](#). *Science*, 358(6370), 1530-1534.
2. Brynjolfsson, E., Thierer, A., & Acemoglu, D. (2024, March). [Navigating the Future of Work: Perspectives on Automation, AI and Economic Prosperity](#). AEI.
3. Brynjolfsson, E. (2022). [The Turing Trap: The promise & peril of human-like artificial intelligence](#). *Daedalus*, 151(2), 272-287.

Optional readings:

4. Turing, A. (1950). [Computing Machinery and Intelligence](#).
5. Bieber, C. (2023). [ChatGPT broke the Turing test — the race is on for new ways to assess AI](#). *Nature*, 619(686-689).
6. McElheran, K., et al. (2023). [AI Adoption in America: Who, What, and Where](#). *National Bureau of Economic Research Working Paper* 31788.
7. Brynjolfsson, E., Li, D., & Raymond, L. (2023, April). [Generative AI at Work](#). *National Bureau of Economic Research Working Paper* 31161.
8. 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.
9. Bubeck, S., et al. (2023). [Sparks of Artificial General Intelligence: Early experiments with GPT-4](#). arXiv preprint arXiv:2303.12712.

April 9: [Eric Schmidt] The Age of AI

Main readings:

1. Schmidt, E. (2023, April). [Does One Model Rule Them All? – Predictions on the Future AI Ecosystem](#).
2. Schmidt, E. (2022). [AI, Great Power Competition & National Security](#). *Daedalus*, 151(2), 288-298.
3. Kissinger, H., Schmidt, E., & Huttenlocher, D. (2023, February 24). [ChatGPT Heralds an Intellectual Revolution](#). *The Wall Street Journal*.
4. Brynjolfsson, E., & Unger, G. (2023, December). [The Macroeconomics of Artificial Intelligence](#). *IMF*.
5. Sutton, R. (2019, March). [The Bitter Lesson](#). [blog post]

Optional readings:

6. Schmidt, E. (2023, April 6). [The Mystery Ahead of Us for Mankind](#). *Special Competitive Studies Project*.
7. Suleyman, M. (2023). Four Features of the Coming Wave (Chapter 7). In *The Coming Wave: Technology, Power, and the Twenty-First Century's Greatest Dilemma*. Crown. (see PDF on Canvas)
8. Kissinger, H. A., Schmidt, E., & Huttenlocher, D. (2021). [The age of AI: and our human future](#). Hachette UK.
9. Schmidt, E. (2023). [Innovation Power: Why Technology Will Define the Future of Geopolitics](#). *Foreign Affairs*.
10. Council of Economic Advisors. (2024). An Economic Framework for Understanding Artificial Intelligence. In *The Economic Report of the President*. (see PDF on Canvas)
11. Metaculus predictions:
 - a. [When will the first general AI system be devised, tested, and publicly announced?](#)
 - b. [When will we have transformative AI?](#)

April 16: [Mira Murati] LLMs and Beyond

Main readings:

1. Video generation models as world simulators, [OpenAI \(2024\)](#)
2. Our approach to AI safety ([OpenAI 2023](#))
3. Moving AI governance forward ([OpenAI 2023](#))
4. How OpenAI is approaching worldwide elections ([OpenAI 2024](#))
5. Tech Accord to Combat Deceptive Use of AI in 2024 Elections ([AI Elections Accord 2024](#))
6. Manning, C. D. (2022). [Human language understanding & reasoning](#). *Daedalus*, 151(2), 127-138.

Optional readings:

7. AI and Catastrophic Risk, [Bengio \(2023\)](#)
8. AGI Safety from First Principles ([Ngo 2020](#))
9. Clarifying "AI Alignment" ([Christiano 2018](#))
10. Unsolved problems in ML Safety ([Hendrycks et al. 2022](#))
11. How do social media feed algorithms affect attitudes and behavior in an election campaign? ([Guess et al. 2023](#))
12. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality, [Dell'Acqua et al. \(2023\)](#)
13. How Generative AI Can Augment Human Creativity. ([Eapen et al. 2023](#))
14. Regulating Transformative Technologies ([Acemoglu et Lensman 2023](#))
15. Wolfram, S. (2023, February 14). [What is ChatGPT doing... And Why Does it Work?](#) [blog post]. (skim).

16. Vaswani, A., et al. (2017). [Attention Is All You Need](#). arXiv preprint arXiv:1706.03762.
17. Joshi, S. (2023). [Everything you need to know about: Scaling Laws in Deep Learning](#). [Blog post].
18. OpenAI. (2020). [Scaling Laws for Neural Language Models](#). arXiv preprint arXiv:2001.08361v1.
19. Thiriet, A. (2023). [Scaling Laws and Emergent Properties](#). [Blog post].
20. Christiano, P., et al. (2017). [Deep Reinforcement Learning from Human Preferences](#). arXiv preprint arXiv:1706.03741.
21. Towards Monosemanticity: Decomposing Language Models With Dictionary Learning ([Bricken et al. 2023](#))
22. How likely is deceptive alignment? ([Hubinger 2022](#))
23. Schaeffer, R., Miranda, B., & Koyejo, S. (2024). [Are emergent abilities of large language models a mirage?](#) *Advances in Neural Information Processing Systems*, 36.
24. Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., et al. (2021). [On the opportunities and risks of foundation models](#). arXiv preprint arXiv:2108.07258. (skim).

April 23: [Reid Hoffman] AI and Humanity

Main readings:

1. Hoffman, R. (2023, January). [Technology Makes Us More Human](#). *The Atlantic*.
2. Hoffman, R. (2023) [Bologna Business School Commencement Speech](#) [blog post]
3. Hoffman, R. (2023). Introduction, pp. 1-25. In *Impromptu*, Dallepedia LLC.
4. Hoffman, R. (2024) Cognitive GPS - draft / do not share (see PDF on Canvas)

Optional readings:

5. Klein, E. (2023, April 16). [The Surprising Thing AI Engineers Will Tell You if You Let Them](#). *The New York Times*.
6. McKinsey. (2023). [Economic Potential of Generative AI](#).
7. Maslej, N., Fattorini, L., Perrault, R., Parli, V., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J. C., Shoham, Y., Wald, R., & Clark, J. (2024, April). [The AI Index 2024 Annual Report](#). AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford, CA.

April 30: [Anima Anandkumar] AI and Creativity

Main readings:

1. Azizzadenesheli, K., Kovachki, N., Li, Z., Liu-Schiaffini, M., Kossaifi, J. and Anandkumar, A., 2024. [Neural operators for accelerating scientific simulations and design](#). *Nature Reviews Physics*, pp.1-9.

2. Song, P., Yang, K., & Anandkumar, A. (2023, November 15). [Towards Large Language Models as Copilots for Theorem Proving in Lean](#). OpenReview.
3. Urban, M., Děchtěrenko, F., Lukavský, J., Hrabalová, V., Svacha, F., Brom, C., & Urban, K. (2024). [ChatGPT improves creative problem-solving performance in university students: An experimental study](#). *Computers & Education*, 105031.
4. Schmidt, E. <https://www.technologyreview.com/2023/07/05/1075865/eric-schmidt-ai-will-transform-science/>
5. Raghu, M., & Schmidt, E. (2020). [A survey of deep learning for scientific discovery](#). arXiv preprint arXiv:2003.11755.
6. Hubert, K. F., Awa, K. N., & Zabelina, D. L. (2024). [The current state of artificial intelligence generative language models is more creative than humans on divergent thinking tasks](#). *Scientific Reports*, 14(3440).

Optional readings:

7. Wang, G., Xie, Y., Jiang, Y., Mandlekar, A., Xiao, C., Zhu, Y., Fan, L. and Anandkumar, A., 2023. [Voyager: An open-ended embodied agent with large language models](#). *arXiv preprint arXiv:2305.16291*.

May 7: [James Manyika] AI and Society

Main readings:

1. Manyika, J., & Spence, M. (2023, October 24). [The Coming AI Economic Revolution: Can Artificial Intelligence Reverse the Productivity Slowdown?](#) *Foreign Affairs*.
2. Ben-Ishai, G., Dean, J., Manyika, J., Porat, R., Varian, H., & Walker, K. (2024, January 24). [AI and the Opportunity for Shared Prosperity: Lessons from the History of Technology and the Economy](#). arXiv preprint arXiv:2401.09718.

Optional readings:

3. [On AI, society and what comes next](#). (2023, December 23). *The Atlantic*.
4. [Dialogues on Technology and Society](#). (2023). YouTube Series.
5. Manyika, J. M. (2022, Spring). [Getting AI Right: Introductory Notes on AI & Society](#). *Daedalus*.

May 14: [Shivon Zilis] Augmenting Humans

Main readings:

1. [Neuralink Show and Tell](#) (video - 2022)
2. Marblestone, A. H., Zamft, B. M., Maguire, Y. G., Shapiro, M. G., Cybulski, T. R., Glaser, J. I., ... Kording, K. P. (2013). [Physical Principles for Scalable Neural Recording](#).
3. Corse, A. (2024, March 23). [Patient of Elon Musk's Neuralink Shows Off New Life With Implant](#). *The Wall Street Journal*.

Optional readings:

4. Musk, E., Neuralink. (2019). [An Integrated Brain-Machine Interface Platform With Thousands of Channels](#). *Journal of Medical Internet Research*, 21(10), e16194. doi: 10.2196/16194
5. Labbé, A. (2023). [Augmented Human: How Generative AI Can Be an Extension of Ourselves](#). *Forbes*.
6. Urban, T. (2017). [Neuralink and the Brain's Magical Future](#). *Wait but Why*.

May 21: [Lawrence Lessig] AI & Democracy

Main readings:

1. Lessig, L. (2021) "The First Amendment Does Not Protect Replicants" (see PDF on Canvas)
2. Lessig, L. (2024). "Protected Democracy" (see PDF on Canvas)

Optional readings:

A bit more pessimistic views:

- Allen, D. and Weyl, E.G., 2024. The Real Dangers of Generative AI. *Journal of Democracy*, 35(1), pp.147-162
- Freedom House, 2023. The repressive power of artificial intelligence. Technical Report. *Freedom House*. <https://freedomhouse.org/report/freedom-net/2023/repressive-power-artificial-intelligence>.
[I suggest using the link rather than the PDF since the publication has interactive maps and videos]

More nuanced overview of gen AI and deliberative democracy:

- Braley, A., Chen L.,, Enríquez J.R.,, Pentland S.,, Reuel A.,, and Tsai L., 2024. Generative AI for Pro-Democracy Platforms. *An MIT Exploration of Generative AI*

Applications:

- Ganguli, D. et al., 2023. Collective constitutional AI: Aligning a language model with public input. *Anthropic*. <https://www.anthropic.com/index/collective-constitutionlai-aligning-a-language-model-with-public-input>.

- The Computational Democracy Project, 2023. Featured Case Studies. Polis Knowledge Base. <https://compdemocracy.org/Case-studies/>

Applied papers:

- Argyle, L.P., Bail, C.A., Busby, E.C., Gubler, J.R., Howe, T., Rytting, C., Sorensen, T. and Wingate, D., 2023. Leveraging AI for democratic discourse: Chat interventions can improve online political conversations at scale. *Proceedings of the National Academy of Sciences*, 120(41), p.e2311627120.
- Motoki, F., Pinho Neto, V. and Rodrigues, V., 2024. More human than human: Measuring ChatGPT political bias. *Public Choice*, 198(1), pp.3-23.

The 2024 election:

- Bueno de Mesquita, E., Canes-Wrone, B., Hall, A.B., Lum K., Martin, G., and Velez, Y. R., 2024. Preparing for Generative AI in the 2024 Election: Recommendations and Best Practices Based on Academic Research. *The University of Chicago Harris School of Public Policy and the Stanford Graduate School of Business*.

May 28: [Daniel Susskind] A World Without Work

Main readings:

1. Susskind, D. (2020). "Part Three: The Response." In *A World Without Work*. Metropolitan Books. (see PDF on Canvas)
2. Russell, S. (2021). Lecture 3: AI and the Economy. In *BBC Reith Lectures*. (see PDF on Canvas)
3. Altman, S. (2021, March 16). [Moore's Law for Everything](#). [blog post]

Optional readings:

1. Susskind, D. (2022). Technological Unemployment. In A. Korinek & J. Bullock (Eds.), *Oxford Handbook of AI Governance*. Oxford University Press. (see PDF on Canvas)
2. Wynroe, et al. (2023). [Literature Review of Transformative Artificial Intelligence Timelines](#). [Blog post].
3. Korinek, A. (2023). [Scenario planning for an A \(G\) I future](#). *F&D, IMF*.
4. Autor, D. (2024, February). [AI Could Actually Help Rebuild the Middle Class](#). *Noema*.
5. Korinek, A., & Stiglitz, J. E. (2018). [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.
6. Acemoglu, D., & Restrepo, P. (2018). [Artificial intelligence, automation, and work](#). In *The Economics of Artificial Intelligence: An Agenda* (pp. 197-236). University of Chicago Press.

7. Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). [GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models](#). arXiv preprint arXiv:2303.10130.
8. Nordhaus, W. D. (2021). [Are we approaching an economic singularity? Information technology and the future of economic growth](#). *American Economic Journal: Macroeconomics*, 13(1), 299-332.
9. Korinek, A., & Suh, D. (2024, March). [Scenarios for the Transition to AGI](#). *National Bureau of Economic Research Working Paper 32255*.