

Resilience in the Age of Artificial Intelligence

Jennifer L. Steele

American University

March 2026

Outline

- 1 What Really Is AI?
- 2 What Will Happen to Jobs?
- 3 Prediction Model Results
- 4 What Does This Mean for All of Us?

Warm-Up

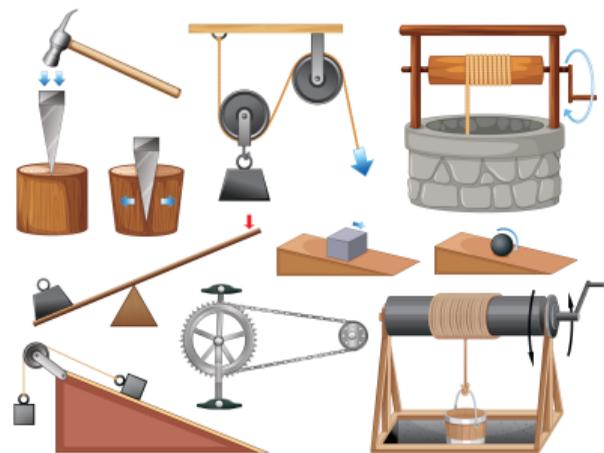


Ask yourself...

- 1 What technological innovation has most surprised you during your lifetime?
- 2 What is your memory of first learning about it?

Tools are part of the human story

- Human success has depended on using tools
- Division of labor underpins all modern economies
- Tools change comparative advantages, affecting economic distribution



www.freepik.com/free-photos-vectors/simple-machines

But AI presents unique dilemmas

Automation of manufacturing jobs increased:

- returns to higher education (Goldin & Katz, 2008)
- economic and political polarization (D. H. Autor & Dorn, 2013)
- all-cause mortality of affected groups age 45–54-(O'Brien et al., 2022)

AI has potential to:

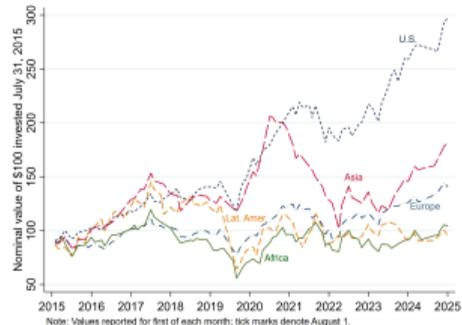
- extend the reach of human expertise (D. H. Autor, 2024)
- accelerate scientific discovery (Hao et al., 2026)
- change the nature of work (D. Autor et al., 2026)



My Wife and My Mother-in-Law, W. E. Hill, 1915

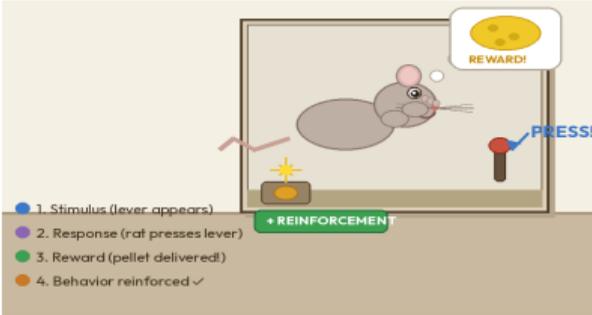
AI is not new

- Balancing an investment portfolio
- Searching Google for information
- Translating a document to another language
- Getting film suggestions from Netflix
- Having Google Maps suggest your route



But is AI "intelligent"?

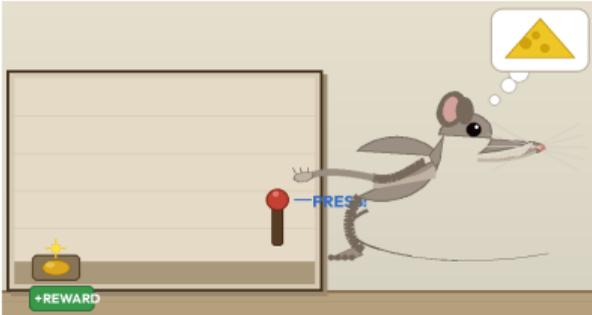
1



2



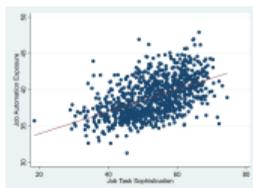
3



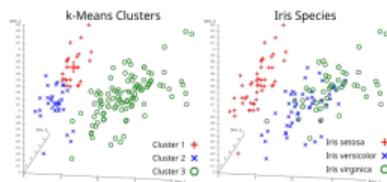
4



How Machines Learn



Author



Chire (Public Domain)



Claude.ai

Supervised Learning

The computer learns from labeled examples to make predictions.

Ex: Predicting the word in Spanish from the word—and neighboring words—in English

Goodfellow, Bengio & Courville (2016) *Deep Learning*

Unsupervised Learning

The computer finds hidden patterns in unlabeled data.

Ex: A streaming service grouping viewers with similar tastes.

Jain, Murty & Flynn (1999) *ACM Computing Surveys*

Reinforcement Learning

A program learns by trial and error, earning rewards for good moves.

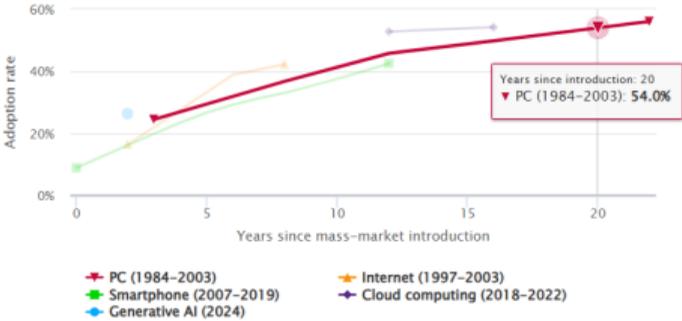
Ex: ChatGPT improved by being rewarded for helpful responses.

Sutton & Barto (2018) *Reinforcement Learning: An Introduction*

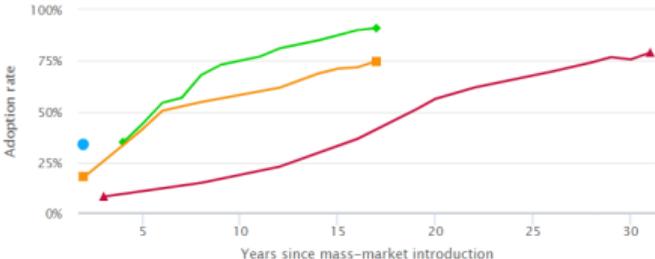
Generative AI adoption has been *fast*

Source: Arnon (2025), citing Bick et al. (2024)
Figure 3. Adoption of Generative AI and Previous Technologies

(a) At Work



(b) Outside of Work



"Attention" to words in context was the breakthrough

- Large language models (LLMs) sounded *human*
- Retrieval augmented generation (RAG) combined trained models with up-to-date info
- Agentic AI can now “remember,” reason through steps, control your computer

Provided proper attribution is provided, Google hereby grants permission to reproduce the tables and figures in this paper solely for use in journalistic or scholarly works.

Attention Is All You Need

Ashish Vaswani¹
Google Brain
avaswani@google.com

Noam Shazeer²
Google Brain
noan@google.com

Niki Parmar²
Google Research
nikip@google.com

Jacob Uszkoreit¹
Google Research
usz@google.com

Llion Jones¹
Google Research
llion@google.com

Aidan N. Gomez¹
University of Toronto
aidan@cs.toronto.edu

Lukas Kaiser²
Google Brain
lukaskaiser@google.com

Iliia Polosukhin¹
iliiia.polosukhin@gmail.com

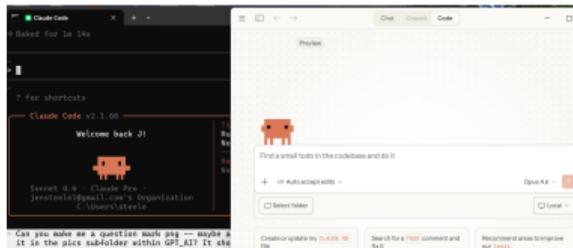
Abstract

The dominant sequence transduction models are based on complex recurrent or convolutional neural networks that include an encoder and a decoder. The best performing models also connect the encoder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train. Our model achieves 28.4 BLEU on the WMT 2014 English-to-German translation task, improving over the existing best results, including ensembles, by over 2 BLEU. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.8 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature. We show that the Transformer generalizes well to other tasks by applying it successfully to English constituency parsing both with large and limited training data.

¹Equal contribution. Listing order is random. Jakob proposed replacing RNNs with self-attention and started the effort to evaluate this idea. Ashish, with Iliia, designed and implemented the first Transformer models and has been crucially involved in every aspect of this work. Noam proposed scaled dot-product attention, multi-head attention and the parameter-free position representation and became the other person involved in nearly every detail. Niki designed, implemented, tuned and evaluated countless model variants in our original codebase and `tensor2tensor`. Llion also experimented with novel model variants, was responsible for our initial codebase, and efficient inference and visualizations. Lukasz and Aidan spent countless long days designing various parts of and implementing `tensor2tensor`, replacing our earlier codebase, greatly improving results and massively accelerating our research.

²Work performed while at Google Brain.
³Work performed while at Google Research.

31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA.



Vibe code in Claude Code. Do forms in CoWork.

Outline

- 1 What Really Is AI?
- 2 What Will Happen to Jobs?**
- 3 Prediction Model Results
- 4 What Does This Mean for All of Us?

Our new research

"Helping People Choose Careers in the Age of AI," (Steele & Cruz, 2026)

- 1 Claude and ChatGPT queries used to estimate automation risk for 923 jobs based on 19,000 tasks and 41 Generalized Work Activities (GWAs)
- 2 Synthesized with projections from five other models
- 3 We find entrepreneurial and social jobs are somewhat less exposed across models

Our analysis compares diverse automation prediction models

Models Under Consideration

Article	How	Measure
Steele & Cruz (2026)	Anthropic & GPT queries	% of job tasks automatable
Eloundou et al. (2024)	GPT & human raters	% of job LLM+ can do twice as fast
Felten et al. (2021)	Crowd-sourced	Job ability automation suitability
Webb (2020)	Text-mining for semantic overlap	AI patent filings & job descriptions
Brynjolfsson & Mitchell (2017)	Crowd-sourced w/ rubric	Task suitability for machine learning (SML)
Frey & Osborne (2017)	Human raters	Abilities that are not social, creative, dexterous

Our empirical measure employs 2025 Claude and OpenAI usage

Using query data from 2025 for Claude (Free, PRO, API) (Appel et al., 2025) and ChatGPT (Chatterji et al., 2025), we define the following:

General Work Activity Exposure defined as:

=**80%** if in the top decile of Claude, Claude API, or OpenAI queries (>7% of queries)

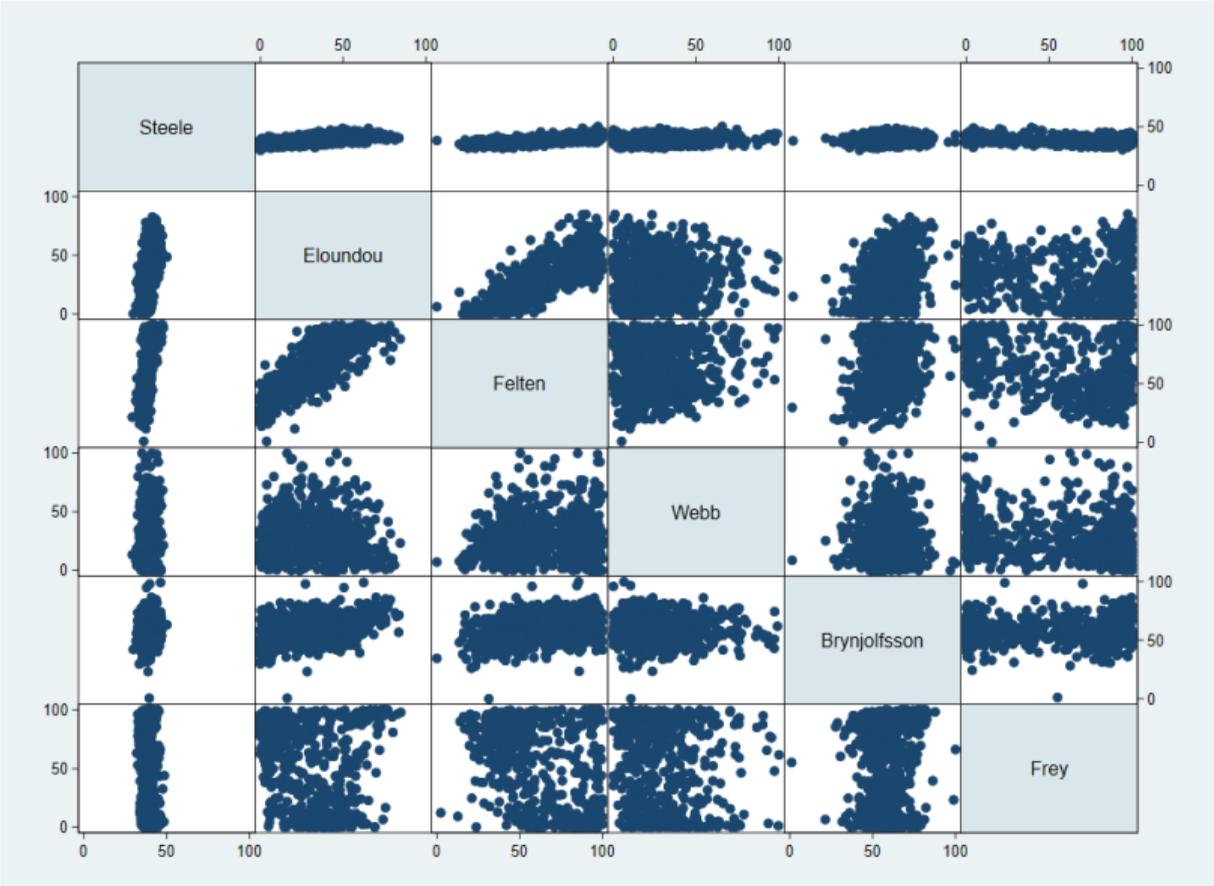
=**45%** if 50th to 90th percentiles (1-7% of queries)

=**10%** if below 50th percentile (<1% of queries)

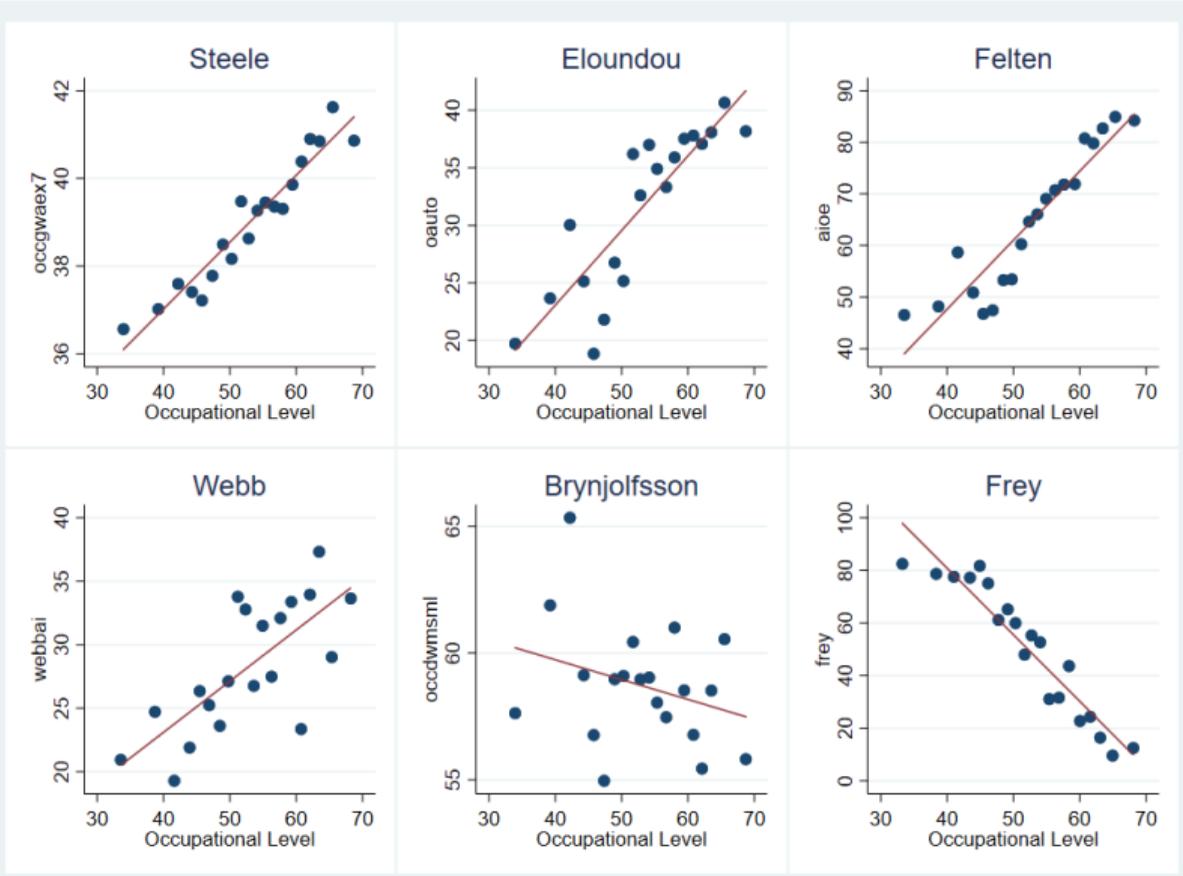
Outline

- 1 What Really Is AI?
- 2 What Will Happen to Jobs?
- 3 Prediction Model Results**
- 4 What Does This Mean for All of Us?

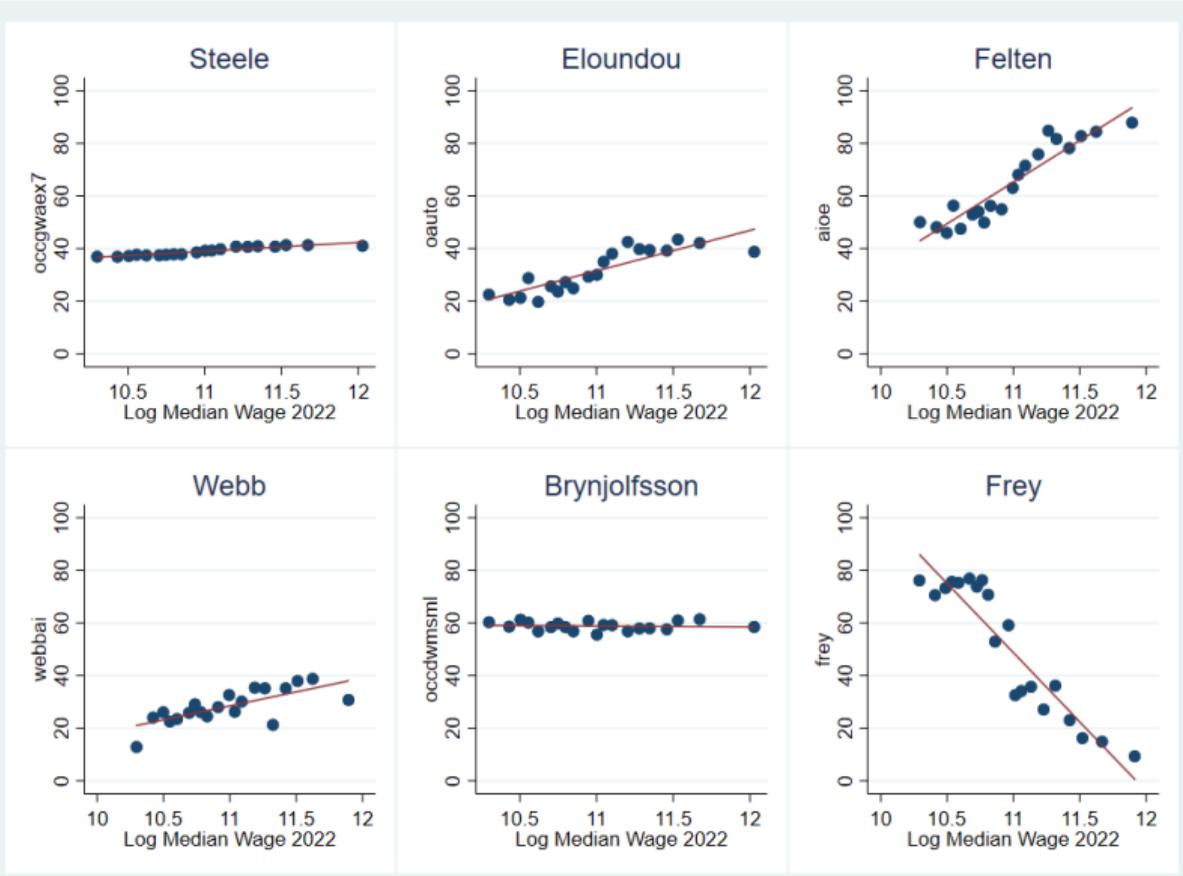
Models yield heterogeneous signals



Models vary in predictions by occupational level



Models vary in predictions by salary



Job category exposure by model

Highest exposure categories and approx % of job tasks exposed, by model

Model	JobZone	RIASEC	Sector
Steele (Queries)	Adv (41.6)	Investigat (41.6)	Compu/Math (43.3) Legal (43.3)
Eloundou (GPT)	Bach (45.9)	Conventional (53)	Office (60.7) Compu/Math (57.6)
Felten (Crowd)	Adv (86.4)	Investigat (85.7)	Compu/Math (97.6) Legal (96.5)
Webb (Patents)	Bach (33.4)	Investigat (38.8)	Compu/Math (45.3) Sciences (45.1)
Brynjolf (Crowd)	Bach (61.6)	Conventional (66)	Office (72) Sales (67.9)
Frey (Theory)	<HS (82.6)	Conventional (67)	Office (84.6) Manufacturing (82.5)

Most-exposed jobs by model

Which jobs show the highest automation exposure, by model?

<i>Steele (Queries)</i> Environmental Economist Mathematician	<i>Eloundou (GPT)</i> Telemarketer Credit Authorizer	<i>Felten (Crowd)</i> Genetic Counselor Financial Examiner
<i>Webb (Patents)</i> Wastewater Treatment Civil Engineer Tech	<i>Brynjolf (Crowd)</i> Mechanical Drafter Mortician	<i>Frey (Theory)</i> Telemarketer Insurance Underwriter

Outline

- 1 What Really Is AI?
- 2 What Will Happen to Jobs?
- 3 Prediction Model Results
- 4 What Does This Mean for All of Us?**

What is certain is uncertainty

In times of rapid transformation, the only constant is...



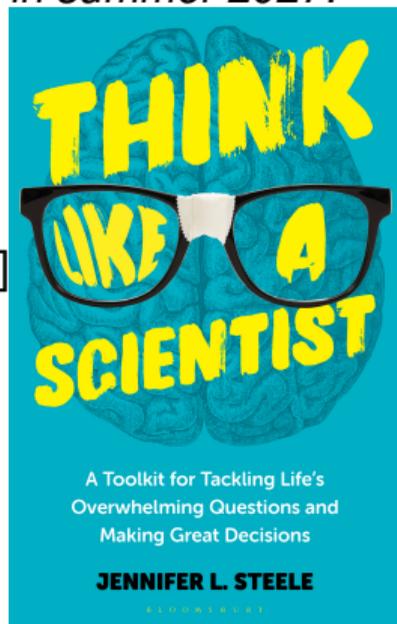
Resilience is the *uber* skill



- Skill demand changes will reshape the economy
- Courage to *learn, pivot, make mistakes* is vital
- Smart bets: entrepreneurship and strategy
- Safe bets: interpersonal and dexterity skills
- Most require EQ: self-regulation and working across difference

Fun links and where to find me

*Coming from Bloomsbury
in summer 2027:*



- **Get started with agents:** Claude CoWork: support.claude.com/en/articles/13345190-get-started-with-cowork (\$20/mo for Pro)
- **Get started with vibe-coding:** Claude Code: code.claude.com/docs/en/quickstart (incl. w/ Pro)
- **Find these slides:** jensteele1.github.io/research/school-to-work
- **Follow me on LinkedIn:** www.linkedin.com/in/jensteeledc
- **Follow me on X:** x.com/jensteeledc

References I

- Appel, R., McCrory, P., Tamkin, A., McCain, M., Neylon, T., & Stern, M. (2025). *The Anthropic Economic Index report: Uneven geographic and enterprise AI adoption* (tech. rep.). <https://assets.anthropic.com/m/218c82b858610fac/original/Economic-Index.pdf>
- Arnon, A. (2025). The projected impact of generative AI on future productivity growth. Retrieved August 9, 2025, from <https://budgetmodel.wharton.upenn.edu/p/2025-09-08-the-projected-impact-of-generative-ai-on-future-productivity-growth/>
- Autor, D., Korinek, A., & Sarin, N. (2026, February). What if labor becomes unnecessary. <https://www.nytimes.com/2026/02/04/opinion/ai-jobs-employment-industry.html>
- Autor, D. H. (2024). Applying AI to rebuild middle class jobs. *NBER Working Paper Series*, w32140. <https://doi.org/10.3386/w32140>
- Autor, D. H., & Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, 103(5), 1553–1597. <https://doi.org/10.1257/aer.103.5.1553>
- Bick, A., Blandin, A., & Deming, D. J. (2024). *The rapid adoption of generative AI*. National Bureau of Economic Research. <https://doi.org/10.3386/w32966>

References II

- Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, 358(6370), 1530–1534. <https://doi.org/10.1126/science.aap8062>
- Chatterji, A., Cunningham, T., Deming, D., Hitzig, Z., Ong, C., Shan, C. Y., & Wadman, K. (2025, September). *How people use ChatGPT* (tech. rep.). National Bureau of Economic Research. Cambridge, MA. <https://doi.org/10.3386/w34255>
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702), 1306–1308. <https://doi.org/10.1126/science.adj0998>
- Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217. <https://doi.org/10.1002/smj.3286>
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280. <https://doi.org/10.1016/j.techfore.2016.08.019>
- Goldin, C., & Katz, L. F. (2008). *The race between education and technology*. Belknap Press.

References III

- Hao, Q., Xu, F., Li, Y., & Evans, J. (2026). Artificial intelligence tools expand scientists' impact but contract science's focus. *Nature*, *649*(8099), 1237–1243.
<https://doi.org/10.1038/s41586-025-09922-y>
- O'Brien, R., Bair, E. F., & Venkataramani, A. S. (2022). Death by robots? Automation and working-age mortality in the United States. *Demography*, *59*(2), 607–628.
<https://doi.org/10.1215/00703370-9774819>
- Steele, J. L., & Cruz, I. (2026). *Helping people choose careers in the age of AI*. American University.
- Webb, M. (2020). The impact of artificial intelligence on the labor market. *SSRN Working Papers*.
<https://doi.org/10.2139/ssrn.3482150>