

Columbia University
School of International and Public Affairs

Our AI Future

DSPC IA7175, Spring 2026

Instructor: **Daniel Björkegren**
dan.bjork@columbia.edu

Meetings: **Wednesdays 1.10-3pm** in IAB 324

Office Hours: A variety of times are available, both virtual and in person. Sign up for a slot using the link on my website (<http://dan.bjorkegren.com>). If none of the listed times work for you, you may email about other potential meeting times.

AI is rewriting the rules of society. This course invites you to understand and shape what comes next. We begin by turning the classroom into a living experiment on how AI could change education, then examine how abundant intelligence could reshape work, governance, and transportation. In a field often dominated by speculation, we will ground our discussions in evidence and theory. Together, we'll explore what institutions are needed for a world transformed by intelligence.

The topics in this course are dynamic, and we will respond to advancements as they occur.

Students will benefit from experience with microeconomics, statistics, and coding, but those are not strict prerequisites.

Our AI Future

Course Components

Class readings will consist of online writing, research papers, and book excerpts. We will ground our discussions based on first principles, and, where possible, data. Given the dynamic nature of these topics, readings are subject to change. Students are encouraged to suggest readings they find thought provoking.

Student discussion is the core of this course. You will be expected to have read any readings, to come to class with questions and observations of your own, and to actively engage with questions and comments. Students are expected to attend all classes unless they have an excused absence, which may include severe medical issues, family emergencies, or significant career-related activities.

Assignments themselves are affected by AI. In the pre-ChatGPT era, I would have asked you to submit 4 memos and a term paper. But AI has expanded what you are capable of creating. So our first task in this class is a meta task: determining what work output would help you learn the most, and be most useful for your careers and impact on the world.

First assignment

Your first assignment, which will be due two days before our second class meeting, is to submit a half page proposal for the remainder of assignments of the course. In our second meeting, we will discuss selected proposals and the class will collectively decide on a single series of assignments that will contribute the most to your learning and demonstrate mastery. We'll use this discussion as a springboard to evaluate the impact of AI on how and what our societies teach.

I anticipate that these proposals will include several regular assignments, such as, for example, writing memos, critiquing memos or arguments, or producing demos to test hypotheses. There are placeholder due dates in the schedule below, though the number and timing of these will be determined during our discussion.

Any technical analysis or complex math should be optional for assignments.

Our AI Future

Update:

Regular assignments. Per the recommendation of the class, students will be asked to complete 4 regular assignments. These will include:

1. *Demo and analysis* (individual or teams of up to 2 people).
 - a. *Demo.* AI makes it possible for people with diverse skillsets to create new things (particularly, software and analysis). In addition, there have been dramatic advances in AI capabilities for creating code, and some are suggesting that those advances may be soon replicated in other forms of knowledge work. Your task is to identify a need and create a minimum viable product of a new product or service, using AI assistance. This need not be technically sophisticated but should demonstrate some new possibility in a domain you care about. You will be asked to send your demo to a relevant classmate a week in advance, and then briefly present it to class.
 - b. *Analysis of classmate's demo.* You will also be asked to constructively analyze the demo of another classmate, using questions such as:
 - i. What is the bigger picture potential of the proposed system?
 - ii. What decision does this system influence?
 - iii. Who would benefit if it works well; who would be harmed if it fails?
 - iv. How might institutions adapt around it; what policies, norms, or institutions currently govern it, or should?
 - c. The demo and the corresponding analysis will be briefly presented in the same class.
2. *Scenario planning* (individual). Students will choose to scenario plan either how (1) their own career, or (2) a particular topic from the course could respond to different scenarios of AI advances.
3. *Short response* (individual). Students will be assigned one topic from the course, and asked to perform additional research and briefly describe their findings to the class.

The deadlines will be spaced throughout the course but students will be asked to complete them in different orders so that each topic is covered.

Final Project (individual or teams of up to 2 people)

In addition to regular assignments, you are expected to produce a final artifact. This may take several forms, such as a product demo, business plan, policy proposal, or research paper. This can build on work you have started for the regular assignments, but should meaningfully advance on it.

In the final weeks of the course each group will have the opportunity to present their artifact. These presentations will allow the rest of the class to learn about your work as well as provide ideas and feedback. A final artifact (typically a written document) will be due after reading period.

Our AI Future

Grades are assembled from these components:

Assignments	50% total (each weighted equally)
Final project	30%
Class participation	20%

Artificial Intelligence

AI tools, including LLMs and interfaces like ChatGPT and GitHub Copilot, can improve productivity (Peng et al. 2023). However, AI tools make mistakes that are difficult to detect unless one understands the underlying principles, and using AI tools prematurely can interfere with learning those principles (Bastani et al. 2025).

You are encouraged to use AI assistance in this course thoughtfully to improve your productivity, to more deeply engage with readings and topics, to help you synthesize knowledge, and to assist you in coming up with ideas.

Our AI Future

Regular Assignment Standards (tentative)

Given that we expect you to use AI tools, we expect that your submissions will meaningfully improve on what an LLM generates when prompted with the assignment. This base LLM response to the assignment description thus represents a benchmark.

Regular assignments be graded by a rubric; we will confirm the rubric after assignments are finalized. Here is a tentative example:

	Excellent	Good	Poor
Engagement (40 points)	Deep engagement with the problems. Substantively improves on the benchmark, by including novel insights, integrating with specialized domain knowledge, or synthesizing in a deeper way. Identifies the deeper questions that arise and attempts to find enduring solutions. (35-40 points)	Good engagement with the problems. Moderate improvement over the benchmark. Attempts to solve deeper questions that arise, but does not go out of its way to do so. (20-34 points)	Does not go deeper than the benchmark, does not engage with all problems, or misses the presence of deeper questions. (19 or below)
Understanding (40 points)	Writing displays a clear understanding of the problem. (35-40 points)	Writing is unclear or does not demonstrate full understanding. (20-34 points)	Writing is unclear or demonstrates substantial misconceptions. (19 or below)
Writing (20 points)	Well written: clear and understandable. (18-20)	Writing is understandable. (13-17)	Does not follow proper structure, includes errors, or poorly written. (12 or below)

Bonus points may be awarded for other insights.

Our AI Future

Policies

Office Hours are a great learning opportunity. Please come to office hours with questions on the material covered in class, comments on the course, or if you want to talk about ideas. Please do not use office hours to talk about grades.

Regrades

Requests for reconsideration of grades are not encouraged, and will be accepted only in writing, with a clear statement of what has been misgraded, within one week of receiving the graded assignment. Please submit your full assignment so grading on all questions can be reconsidered.

Academic Integrity

The School of International & Public Affairs does not tolerate cheating or plagiarism in any form. Students who violate the Code of Academic & Professional Conduct will be subject to the Dean's Disciplinary Procedures. Familiarize yourself with proper citation methods using resources available online: <http://bulletin.columbia.edu/sipa/academic-policies/>. Report violations to the Associate Dean for Student Affairs.

Disability Statement

SIPA is committed to providing reasonable accommodations for students registered with Columbia University's Disability Services (DS). Students with DS-certified accommodation letters should discuss their accommodations with their instructor. Contact SIPA's DS liaison for additional questions: disability@sipa.columbia.edu or 212-854-8690.

Our AI Future

Class Schedule

This schedule is subject to change based on your interests and advances in the space. Readings for each class are shown in indent, with required readings starred (*).

Class 1. January 21.

Primer on Artificial Intelligence

What is artificial intelligence? How does it relate to statistics?

Explainers:

<https://www.understandingai.org/p/large-language-models-explained-with>

<https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/>

LLM Visualizer <https://bbycroft.net/llm>

Videos

3Blue1Brown neural network explainer

https://www.youtube.com/playlist?list=PLZHQObOWTQDNU6R1_67000Dx_ZCJB-3pi

Karpathy's video "[Let's build GPT from scratch](#)"

Neural networks from scratch <https://nnfs.io>

First Assignment Due (by 11am January 26)

Proposal for assignments for this course, to spur class discussion

Class 2. January 28.

Education

What are the opportunities and risks of AI in education? What does that mean for this course, your career—and the world?

* Meeks, F., Agrawal, A., Tana, E., 2024. *How to go from -1 to 0*. South Park Commons. <https://blog.southparkcommons.com/how-to-go-from-minus-1-to-0/> (Read specifically the contrast between 'demo' and 'memo')

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakçı, Ö., Mariman, R., 2025. *Generative AI without guardrails can harm learning: Evidence from high school mathematics*. Proceedings of the National Academy of Sciences 122, e2422633122. <https://doi.org/10.1073/pnas.2422633122>

Muralidharan, K., Singh, A., Ganimian, A.J., 2019. *Disrupting Education? Experimental Evidence on Technology-Aided Instruction in India*. American Economic Review 109, 1426–1460. <https://doi.org/10.1257/aer.20171112>

Our AI Future

Class 3. February 4.

Capabilities

What is AI capable of? How might its capabilities evolve?

Scaling

* Kaplan, J., McCandlish, S., Henighan, T., Brown, T.B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., Amodei, D., 2020. *Scaling Laws for Neural Language Models*. <https://doi.org/10.48550/arXiv.2001.08361>
(Focus on understanding Figures 1 and 2.)

* Sutton, R., 2019. *The Bitter Lesson*.
<http://www.incompleteideas.net/InIdeas/BitterLesson.html>

* Aschenbrenner, Leopold. 2024. *Situational Awareness*.
<https://situational-awareness.ai>
- *Introduction and Part 1*

Kokotajlo, D., Alexander, S., Larsen, T., Lifland, E., and Dean, R. 2025. *AI 2027*.
<https://ai-2027.com>

[Jack Clark's GPT 2 retrospective](#)

Mitchell, M., Krakauer, D.C., 2023. *The Debate Over Understanding in AI's Large Language Models*. <https://doi.org/10.48550/arXiv.2210.13966>

Stanford HAI, 2024. *AI Index Report*. URL <https://aiindex.stanford.edu/report/>

Narayanan, A., Kapoor, S., 2024. *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference*. Princeton University Press.

□ **Possibly: Assignment Due** (by 12pm February 9)

Our AI Future

Class 4. February 11.

Precedents

Do technology shifts lead to institution shifts?

* Jovanovic, Boyan, and Peter L. Rousseau. 2005. "Chapter 18 - General Purpose Technologies." In *Handbook of Economic Growth*, eds. Philippe Aghion and Steven N. Durlauf. Elsevier, 1181–1224. doi:[10.1016/S1574-0684\(05\)01018-X](https://doi.org/10.1016/S1574-0684(05)01018-X).

Rosenberg, N., Trajtenberg, M., 2004. *A General-Purpose Technology at Work: The Corliss Steam Engine in the Late-Nineteenth-Century United States*. The Journal of Economic History 64, 61–99. <https://doi.org/10.1017/S0022050704002608>

Acemoglu, D., Johnson, S., 2023. *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*, 1st edition. ed. PublicAffairs, New York.

□ **Possibly: Assignment Due** (by 12pm February 16)

Class 5. February 18.

Work

How might AI affect work? What will humans do?

* Ford, H., 1939. *Machines as Ministers to Man*. The New York Times 59, 89.

Impacts on existing jobs

* 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*. <https://doi.org/10.48550/arXiv.2303.10130>

Brynjolfsson, E., Chandar, B., Chen, R., 2025. *Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence*. Stanford Digital Economy Lab.

Smith, M., 2026. *Canaries, Interest Rates, and Timing: More on the Recent Drivers of Employment Changes for Young Workers*. Stanford Digital Economy Lab. URL <https://digitaleconomy.stanford.edu/news/canaries-interest-rates-and-timing-a-more-on-recent-drivers-of-employment-changes-for-young-workers/>

Brynjolfsson, E., Li, D., Raymond, L.R., 2023. *Generative AI at Work*. Working Paper Series. <https://doi.org/10.3386/w31161>

Peng, S., Kalliamvakou, E., Cihon, P., Demirer, M., 2023. *The Impact of AI on Developer Productivity: Evidence from GitHub Copilot*. <https://doi.org/10.48550/arXiv.2302.06590>

Our AI Future

Work and meaning

Hussam, R., Kelley, E.M., Lane, G., Zahra, F., 2022. *The Psychosocial Value of Employment: Evidence from a Refugee Camp*. *American Economic Review* 112, 3694–3724. <https://doi.org/10.1257/aer.20211616>

Universal Basic Income

Bartik, A.W., Rhodes, E., Broockman, D.E., Krause, P.K., Miller, S., Vivalt, E., 2024. *The Impact of Unconditional Cash Transfers on Consumption and Household Balance Sheets: Experimental Evidence from Two US States*. Working Paper Series. <https://doi.org/10.3386/w32784>

Miller, S., Rhodes, E., Bartik, A.W., Broockman, D.E., Krause, P.K., Vivalt, E., 2024. *Does Income Affect Health? Evidence from a Randomized Controlled Trial of a Guaranteed Income*. Working Paper Series. <https://doi.org/10.3386/w32711>

Vivalt, E., Rhodes, E., Bartik, A.W., Broockman, D.E., Miller, S., 2024. *The Employment Effects of a Guaranteed Income: Experimental Evidence from Two U.S. States*. Working Paper Series. <https://doi.org/10.3386/w32719>

Crosta, T., Karlan, D., Ong, F., Rüschenpöhler, J., Udry, C.R., 2024. *Unconditional Cash Transfers: A Bayesian Meta-Analysis of Randomized Evaluations in Low and Middle Income Countries*. Working Paper Series. <https://doi.org/10.3386/w32779>

□ **Possibly: Assignment Due** (by 12pm February 23)

Our AI Future

Class 6. February 25.

Fakes, Signaling, and Safeguards

How might AI affect the ability to signal? What new institutions might be needed?

* Gans, J.S., 2024. *How will Generative AI impact communication?* Economics Letters 242, 111872.

* Costello, T.H., Pennycook, G., Rand, D.G., 2024. *Durably reducing conspiracy beliefs through dialogues with AI.* Science 385, eadq1814.
<https://doi.org/10.1126/science.adq1814>

Manipulation and brittleness

Eykholt, K., Evtimov, I., Fernandes, E., Li, B., Rahmati, A., Xiao, C., Prakash, A., Kohno, T., Song, D., 2018. *Robust Physical-World Attacks on Deep Learning Visual Classification.* Presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1625–1634.

Björkegren, D., Blumenstock, J.E., Knight, S., 2024. *Manipulation-Robust Prediction.*

Possibly: Assignment Due (by 12pm March 2)

Project Proposal Due (by 8pm March 2)

Class 7. March 4.

Emerging Markets

How do opportunities and risks of AI differ among low-income populations?

* Björkegren, D., 2023. *Artificial Intelligence for the Poor.* Foreign Affairs.

Machine Learning

Aiken, E., Bellue, S., Karlan, D., Udry, C., Blumenstock, J.E., 2022. *Machine learning and phone data can improve targeting of humanitarian aid.* Nature 603, 864–870. <https://doi.org/10.1038/s41586-022-04484-9>

Björkegren, D., Grissen, D., 2018. *The Potential of Digital Credit to Bank the Poor.* American Economic Association Papers and Proceedings.

Our AI Future

Artificial Intelligence

Otis, N., Clarke, R., Delecourt, S., Holtz, D., Koning, R., 2024. *The Uneven Impact of Generative AI on Entrepreneurial Performance*.

<https://doi.org/10.2139/ssrn.4671369>

* Björkegren, D., Choi, J.H., Budihal, D., Sobhani, D., Garrod, O., Atherton, P., 2025. Could AI Leapfrog the Web? Evidence from Teachers in Sierra Leone.

<https://doi.org/10.48550/arXiv.2502.12397>

McPeak, G., Sautmann, A., George, O., Hallal, A., Simal, E.A., Schwartz, A.L., Abaluck, J., Ravi, N., Pless, R., 2024. *An LLM's Medical Testing Recommendations in a Nigerian Clinic: Potential and Limits of Prompt Engineering for Clinical Decision Support*, in: 2024 IEEE 12th International Conference on Healthcare Informatics (ICHI). Presented at the 2024 IEEE 12th International Conference on Healthcare Informatics (ICHI), pp. 586–591.

<https://doi.org/10.1109/ICHI61247.2024.00094>

Class 8. March 11.

Alignment

What values should guide automated systems?

* Anthropic. *Core Views on AI Safety: When, Why, What, and How*.

<https://www.anthropic.com/news/core-views-on-ai-safety>

* Kleinberg, J., Ludwig, J., Mullainathan, S., Sunstein, C.R., 2020. *Algorithms as discrimination detectors*. PNAS. <https://doi.org/10.1073/pnas.1912790117>

Kasy, M., 2023. *The political economy of AI: Towards democratic control of the means of prediction*.

Abebe, R., Kasy, M., 2021. *The Means of Prediction*. Boston Review.

No class. March 18. Spring Recess.

□ **Possibly: Assignment Due** (by 12pm March 23)

Our AI Future

Class 9. March 25.

Teaching the machines

Who will teach the machines? How should they be paid?

Returns to data: scaling laws revisited

* Gans, J.S., 2024. *Copyright policy options for generative artificial intelligence*. National Bureau of Economic Research.

* Ghorbani, A., Zou, J., 2019. *Data Shapley: Equitable Valuation of Data for Machine Learning*, in: Proceedings of the 36th International Conference on Machine Learning. Presented at the International Conference on Machine Learning, PMLR, pp. 2242–2251.

□ **Possibly: Assignment Due** (by 12pm March 30)

Class 10. April 1.

Digital governance

Are new forms of governance possible—or necessary?

* *The Digitalist Papers*. <https://www.digitalistpapers.com> (choose two to read)

Weyl, G., Tang, A., 2024. *Plurality: The Future of Collaborative Technology and Democracy*.

Pol.is

<https://communitynotes.x.com/guide/en/about/introduction>

No class. April 8.

□ **Possibly: Assignment Due** (by 12pm April 13)

Our AI Future

Class 11. April 15.

Markets and planning

How can data and algorithms reshape markets?

* Hayek, F.A., 1945. *The Use of Knowledge in Society*. The American Economic Review 35, 519–530.

Brynjolfsson, E., Hitzig, Z., 2025. *AI's Use of Knowledge in Society* (No. c15303). National Bureau of Economic Research.

* Morozov, E., 2019. *Digital Socialism?* New Left Rev 33–67.

Shahidi, P., Rusak, G., Manning, B.S., Fradkin, A., Horton, J.J., 2025. *The Coasean Singularity? Demand, Supply, and Market Design with AI Agents* (No. c15309). National Bureau of Economic Research.

Farronato, C., Fradkin, A., Larsen, B.J., Brynjolfsson, E., 2024. *Consumer Protection in an Online World: An Analysis of Occupational Licensing*. American Economic Journal: Applied Economics 16, 549–79. <https://doi.org/10.1257/app.20210716>

Luca, M., 2016. *Reviews, Reputation, and Revenue: The Case of Yelp.Com*. <https://doi.org/10.2139/ssrn.1928601>

Class 12. April 22.

Project presentations

To be announced

Class 13. April 29.

Project presentations

To be announced

Class 14. Makeup class.

To be announced

□ **Final Project Due** (by 8pm May 8)

Version 14 April 2026