Automation and AI

Jinpeng Shi

October 19, 2026

Lecture Structure

- In the next three lecture, we will cover topics including:
  - AI for Science;
  - The Fundamentals of AI;
  - Effect of Automation;
  - Early Impacts of AI.
- We first discuss our thoughts and feelings on how AI could shape future scientific discovery, then we understand the structure of LLMs and AI agents, we conclude with latest empirical work related to this issue.

1. AI for Science

Motivation

- Motivating discussion: AI for Science

We divide today's discussion into three parts:

- AI for scientific discovery;
- The future of academic training, e.g., PhD programs;
- AI on peer-reviewed publication system.

Q1. How will AI shape the future of scientific discovery?

AI in Hypothesis Generation

- Suppose an AI system proposes a new cancer treatment that no human scientist had previously considered. A human laboratory then verifies it. Who deserves credit for the discovery?
  - Google's Co AI Scientist, Nature 2026
  - Co Scientist from DeepMind

AI in Hypothesis Generation

- Suppose an AI system proposes a new cancer treatment that no human scientist had previously considered. A human laboratory then verifies it. Who deserves credit for the discovery?
  - Google's Co AI Scientist, Nature 2026
  - Co Scientist from DeepMind

So what counts for a scientific discovery?

- Does scientific creativity require human understanding, or is producing a genuinely novel, testable hypothesis sufficient?

AI in Hypothesis Generation Continued

- Q1: If everyone has access to the same AI, will scientists generate more diverse ideas or increasingly pursue similar ones?

AI in Hypothesis Generation Continued

- Q1: If everyone has access to the same AI, will scientists generate more diverse ideas or increasingly pursue similar ones?
- Q2: Are lower ranked hypotheses really less worthy for trying?

AI in Hypothesis Generation Continued

- Q1: If everyone has access to the same AI, will scientists generate more diverse ideas or increasingly pursue similar ones?
- Q2: Are lower ranked hypotheses really less worthy for trying?
- Q3a: What if advantaged/senior researchers always get (earlier) access to frontier models and computing powers?
- Q3b: Will they pursue ideas fall within their research agenda or push the limit of science?

AI in Data Simulation

MatrAIx created 8.3 billion synthetic human personas, approximately matching the scale of the world's population.

- MatrAIx Simulated Agents
- Playground

With such AI generated synthetic data, can they replace actual observations?

AI in Empirical Analysis

An AI agent today can clean the data, choose an econometric model, run regressions, and produce a publication-ready paper.

- Autoresearch
- Which tasks would you be comfortable delegating to AI?

AI in Empirical Analysis

An AI agent today can clean the data, choose an econometric model, run regressions, and produce a publication-ready paper.

- Autoresearch
- Which tasks would you be comfortable delegating to AI?
- Should AI be allowed to choose the identification strategy?
- With a regulated agentic workflow: AI on data cleaning and running regressions under detailed instruction, do we as a research still need to check the code they generated?

Who Validates AI Generated Results?

OpenAI released a math repo on github contains over 700 math manuscripts:

- OpenAI Math
- Who evaluates their results?
- Who gets the credit after a successful validation?
- What if it turns out some of the them were actually false results (40% of the manuscripts do not have lean proof)?
- What about time wasted in understanding AI generated manuscripts?

AI Companies and Scientists Misalignment

- What are your thoughts on Terrance Tao's open letter and Leiden declaration?

AI Companies and Scientists Misalignment

- What are your thoughts on Terrance Tao's open letter and Leiden declaration?
- Should publicly funded AI systems prioritize scientific problems that are commercially unprofitable?

AI Companies and Scientists Misalignment

- What are your thoughts on Terrance Tao's open letter and Leiden declaration?
- Should publicly funded AI systems prioritize scientific problems that are commercially unprofitable?
- Do you think AI will favor questions with more existing data or favor the questions that are less developed today?
- Which type of questions align more closely with scientific pursuits?

Q2. How will AI change the future of academic training?

AI as an Academic Advisor

Imagine choosing between two graduate programs with the same access to AI research assistants:

- Program A: A prestigious human advisor, but limited meeting time.
- Program B: A less prominent human advisor, but more time for one-on-one meetings.

Which one would do you choose? Why?

AI in PhD Education

Many people argued that PhD training should evolve to make the most of AI (Nature, 2025):

- So, do we teach PhD students the same materials as before?
- If not, what should we teach them today and in the near future?
- If we don't teach them the knowledge that we are using today, how can they evaluate the validity of AI generated outputs?

AI in PhD Education

Many people argued that PhD training should evolve to make the most of AI (Nature, 2025):

- So, do we teach PhD students the same materials as before?
- If not, what should we teach them today and in the near future?
- If we don't teach them the knowledge that we are using today, how can they evaluate the validity of AI generated outputs?
- More generally, how can they develop independent scientific judgment if AI generates most of their research ideas?

Pace, Outcome, and Expectation

If AI allows a PhD student to complete research twice as fast, should the PhD take half as long?

- What happens to the length of a PhD?
  - It becomes substantially shorter;
  - It stays similar, but expectations increase;
  - It becomes longer as the research frontier advances;
  - Traditional PhD training becomes less necessary

AI on Academic Inequality

Can AI replace intellectual mentorship?

- Would access to advanced AI tools make elite universities more or less important?
- Would advance of AI make academia more like a club or a decentralized system?

AI on Academic Inequality Continued

What about financial support for PhD students/Postdocs?

- Will PhD programs disappear or increase?
- Isn't that a return of elite education (and the end of merit-based system), which only advantaged families can support their children earning the Doctoral degree?
- Will AI democratize access to scientific careers, or make competition for a small number of faculty positions even more intense?

Q3. What’s left for peer-reviewed publication system under AI?

AI on Paper Polishing

Do you think the following scenarios are acceptable?

- Grammar and language polishing;
  - Is disclosure of AI usage necessary in this case?

AI on Paper Polishing

Do you think the following scenarios are acceptable?

- Grammar and language polishing;
  - Is disclosure of AI usage necessary in this case?
- Rewriting an entire manuscript;
  - Who is responsible for the arguments?

AI on Paper Polishing

Do you think the following scenarios are acceptable?

- Grammar and language polishing;
  - Is disclosure of AI usage necessary in this case?
- Rewriting an entire manuscript;
  - Who is responsible for the arguments?
- Generating literature reviews and results interpretations.
  - When does assistance become authorship?

AI as Reviewers

If AI can write a paper and AI can review it, what role remains for human judgment?

- Would you trust a paper reviewed by three AI systems more than one reviewed by three human experts?

AI as Reviewers

If AI can write a paper and AI can review it, what role remains for human judgment?

- Would you trust a paper reviewed by three AI systems more than one reviewed by three human experts?
- Should journals disclose whether a manuscript was reviewed using AI?

AI as Reviewers

If AI can write a paper and AI can review it, what role remains for human judgment?

- Would you trust a paper reviewed by three AI systems more than one reviewed by three human experts?
- Should journals disclose whether a manuscript was reviewed using AI?
- If AI reduces the cost of reviewing, should journals publish more papers or impose stricter standards?

Food for Thoughts

- What kind of research should we fund in the future?
- If less grant is made to a particular type of research, should university change the standard for tenure?
- Should universities continue evaluating faculty using publication counts when AI dramatically reduces the cost of producing manuscripts?
- Should universities raise the bar for publication outcomes?

2. A Gentle Introduction to Fundamentals of AI

What is Artificial Intelligence (AI)?

John McCarthy (2007):

- "*Artificial intelligence is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable*."

Large Language Models (LLMs)

Language AI:

- refers to a subfield of AI that focuses on developing technologies capable of understanding, processing, and generating human language.
- You may have heard its other names: e.g., Natural Language Processing (NLP), etc.

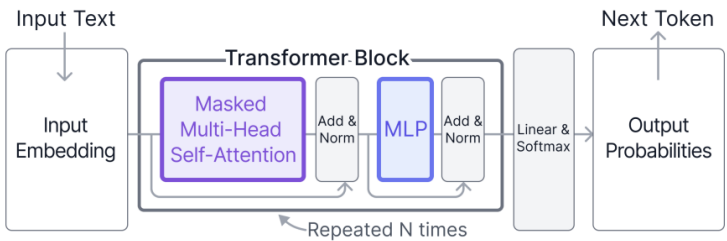
Large Language Models:

- is a specific model class within Language AI.
- LLMs have different models: BERT, GPT, etc.

Transformer

Transformer is introduced in a landmark paper:

- “Attention Is All You Need”, Vaswani et al. (2017).
- Today, it becomes the foundational model for many LLMs we use, like GPT.



Transformer, Neural Network, Deep Learning and Else

- A visual explanation: Transformer Explainer
- Transformer is a type of neural network;
- Neural network is a machine learning method developed around 1950s;
- It hasn't drawn much attention until the recent revival of deep neural network (deep learning);
- Deep neural network essentially is a multi-layer neural network.

More Jargon

- People tend to use AI, deep learning, machine learning and statistical learning interchangeably;
- Statistical learning focus on inferences, interpretability of statistical models;
- Where machine learning is more like a black box machine that are used for model prediction;
- Deep learning is a subset of machine learning method, essentially a multi-layer neural network, makes the model less interpretable;
- AI today refers to a broad area to mimic human intelligence.

Transformer Breakdown: Embedding

Suppose we have the following input:

- This year Nobel Prize in Economics is _____
- First, the input is split into **tokens**: words, subwords, character combinations;
- Each token is then converted into token **embeddings**: a numerical (vector) representation;
- Position encoding is also added to preserve the order of tokens.

Tokenization



More on Tokenization

- Token ID is matched to a vector in a learned matrix (all possible tokens);
- So, in the embedding layer, words (or pixels, audio units) are “translated” into unique vectors;
- Transformer is essentially a complicated matrix multiplication method.
- Note that the learned matrix is trained during **pre-training** process;
- Where AI labs train the model based on enormous information they crawled from the website.

Attention: Context Matters

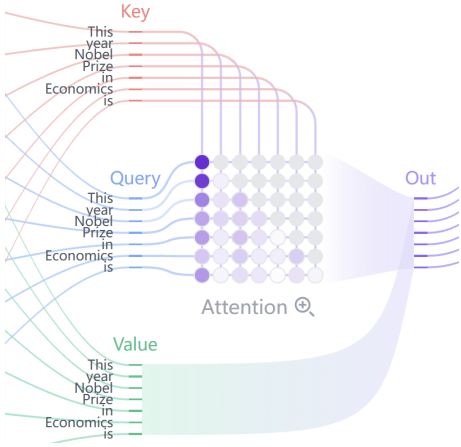
- Then the embeddings are passed through multiple layers of transformer blocks;
- It is a repeated process until we reach the final layer.
- We know that words in different contexts may have different meanings;
- To understand the specific context, we want words/tokens “**talk**” to each other;
- In the attention layer: transformer is effectively calculating the matrix product to update the “value” of embedding vectors.

More on Attention

Common technical terms:

- Query matrix: like a search engine to search for relevant words/tokens;
- Key matrix: like a self-identity, what information this word/token has to pass to other words;
- Attention score: a matrix product of the above information;
- Value matrix: using the information we get from the previous step to update the input embedding $\Rightarrow$ The true semantic meaning.
- Masking: You only give information that comes before the word, not after;

Attention Matrices

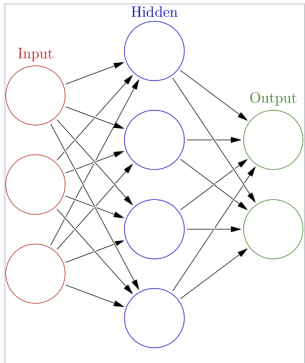


Multi-head Attention

- So far, what we have talked about is a single head attention;
- In practice, we use multi-head attention in parallel.
- During the training process, we need to give a starting number to train the data, so with different starting number, the model may learning different information from the data;
- Thus, we do multi-head attention to capture all possible information;
- And during inference process, each head attention layer may provide different information to update the embedding vectors.

Multi-layer Perceptron (MLP)

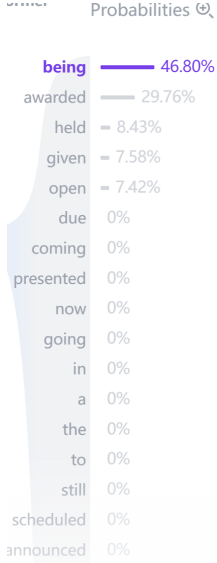
- Now, each token has successfully talked to each other, and we have a new vector that contains context information;
- And we want to process that vector to help us predict the next word/token
- Here the neural network comes in:



Output Layer

- In the last layer, we have a new vector which represents the possible next words based on the input;
- A softmax function transforms that vector into a probability distribution;
- Based on the probabilities of each possible word, the model generate the next word;
- Then this entire process is repeated to generate word sequence following the input.

Output Probabilities



GPT History

GPT: Generative Pre-Trained Transformer

- GPT(Radford et al, 2018): unsupervised learning (pre-trained model) + fine-tuning (labeled data for specific task).
- GPT2 (Radford et al, 2019): zero shot setting is sufficient.
- GPT3 (Brown et al, 2020): NLP is few shot learning. (prompting + in task examples)
- GPT4 (Open AI, 2023): Include image processing. (Multi-modal model)

A Spanish Learning Allegory

To understand zero/few shot training: **Suppose you want to learn Spanish**

- Traditionally, you sign up for courses, study grammar, read articles, watch videos, and talk to native speakers.

Now, what GPT model does is

- Feed you 10,000 books, articles, online posts etc., hope you will discover the pattern of Spanish yourself.
- Then prompt you with specific task, like academic writing, and provide you templates, hope you will learn Spanish academic writing by time.

Why can a chatbot become useful?

- Pre-training: typically train on large data set without specify specific use of task:
  - classification, coding, image generation, etc.
- Post-training: You can treat it as telling the model how you prefer it generating responses.
- Prompt: Give the model your context and instruction.
- Retrieval and Tool calling: gathering external evidence and calling the computational tools.

Hallucination

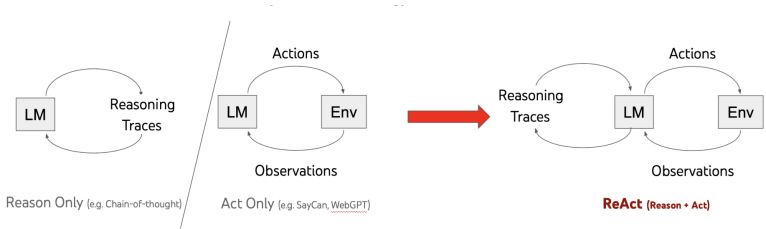
Why a model can make mistake?

- Plausible language generated by LLM is not the same as verified truth:
  - It may generate things that never exist;
  - It may be based on outdated knowledge and information;
  - It may be subject to bias presented in the historic data or AI labs.
- A small error may compound exponentially for long tasks;
- We need external verification: retrieval and/or human oversight.

Reasoning: Chain of Thought

- Chain of thought: shows the intermediate steps for solving the problem;
- It is not guaranteed that it reveals the internal processing.

ReAct: reasoning+acting



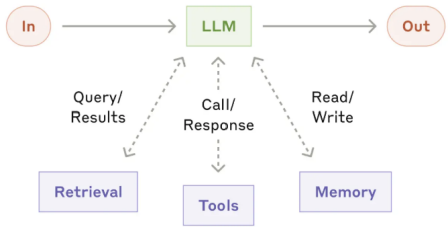
The Agentic AI

Workflow vs Agent:

- Workflow: are systems where LLMs and tools are orchestrated through predefined code paths.
- Tools: calculators, Python (packages), APIs, files, etc.
- Agent: LLMs dynamically direct their own processes and tool usage, maintaining control over how they accomplish tasks.

The RAG

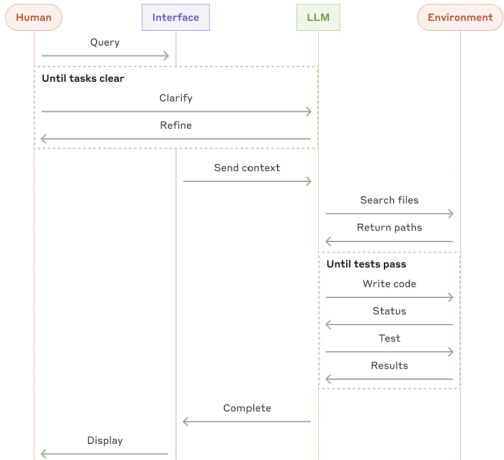
Retrieval augmented generation (RAG):



The Agent Loop

- Agents begin their work with either a command from, or interactive discussion with, the human user.
- Once the task is clear, agents plan and operate independently, potentially returning to the human for further information or judgment.
  - Human in the loop.
- During execution, it's crucial for the agents to gain "ground truth" from the environment at each step (such as tool call results or code execution) to assess its progress.
- Agents can then pause for human feedback at checkpoints or when encountering blockers.
- The task often terminates upon completion, but it's also common to include stopping conditions (such as a maximum number of iterations) to maintain control.

Visualizing the Agent Loop



Harness, memory, skills, MCP

You may have encountered some jargon in agentic AI:

- MCP: the protocol to connect AI applications to tools and data;
- Thus, the tool permissions and untrusted data need safeguards!
- Harness: manage context, tool loop, stopping conditions;
  - You may need a stopping condition, otherwise your tokens may be burnt on some error loops without stopping.
- Skills: you can treat it as a complicated prompt;
- Memory: saved or retrieved context;
- Runtime/Sandbox: where tools execute and code run.

Demo

- AI concept demo

3. The Impact of Automation

Automation: Definition

According Merriam Webster:

- Automatically controlled operation of an apparatus, process, or system by mechanical or electronic devices that take the place of human labor.

Acemoglu and Restrepo (2019):

- Automation corresponds to the development and adoption of new technologies that enable capital to be substituted for labor in a range of tasks.

Past and Future Waves of Automation

- 1st wave: Industrial revolution automated tasks performed by artisans, e.g., spinning and weaving.
- 2nd wave: Mechanization of agriculture for reapers and harvesters.
- 3rd wave: Industrial robots and machinery disrupt jobs of production workers.
- Ongoing wave: Specialized software and artificial intelligence perform tasks in accounting, sales, logistics, trading, and managerial occupations.

The Central Debate

The implications of technological change for employment and wages:

- Some see the ongoing process of automation (robots, AI) as the harbinger of widespread joblessness;
- Others see the current automation, like previous waves of technologies, will ultimately increase labor demand.

Q: How will automation shape future employment and wages?

Task-based Framework in A Nutshell

Consider a production function $F(K, L)$ takes capital and labor as input,

- New technologies **increase** productivity of both inputs;
- Also impact the allocation of tasks to these inputs;
- what tasks are produced by capital or labor (or the shares of the two) under new technologies?

$\Rightarrow$ Changes in task content of production affect **labor demand**.

The Implications of Technology

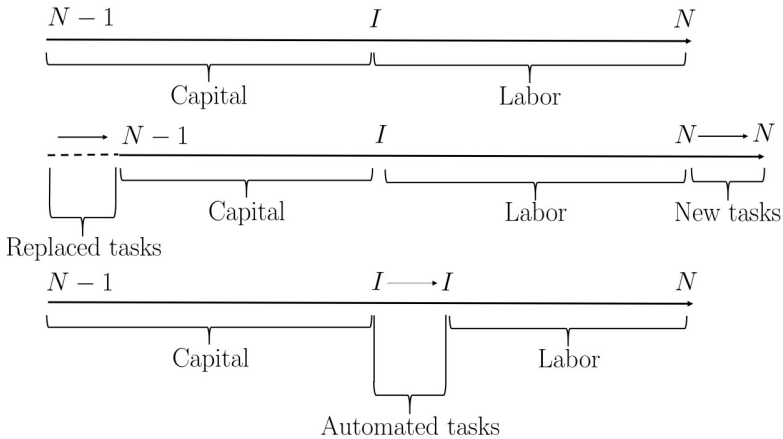
- Displacement effect: capital takes over tasks previously performed by labor.
  - Automation **reduces** the labor share of value added.
- Productivity effect: affect the labor demand in non-automated tasks.
  - Automation **increases** productivity.
- Reinstatement effect: create new tasks which labor has a comparative advantage.
  - Automation change task content **in favor of** labor.

Task-based Framework: Set up

Consider a single sector economy:

- Production combines the output of a range of tasks (lie in $N-1$ and N).
- Tasks can be produced using capital or labor.
- Consider a threshold $I \in [N-1, N]$,
  - Tasks above I can only be produced by labor;
  - Tasks below can be automated, produced with capital as well as labor.
- Assume cost-minimizing firms and immediate adoption of new technologies.
  - Automation: an increase in I;
  - New tasks: an increase in N;

Allocation of Labor



Technology and Labor Demand: Set up

Q: What's the impact of technological change on labor demand?

Define wage bill as WL,

- W: labor market wage. In equilibrium, it equals to workers' value added.
- L: labor share.

We consider two types of technology impact on labor demand:

- Effect of **automation**;
- Effect of **new tasks**;

Technology and Labor Demand: Automation Implication

Effect of **automation** on labor demand $=$ Productivity effect $+$ Displacement effect

- Productivity effect:
  - Automation increases value added;
  - Raise demand for labor from non-automated tasks.
- Displacement effect:
  - Displace labor from tasks previously assigned to;
  - Reduce labor share.

Thus, so-so technologies only generate small productivity effect, **not enough** to offset labor share decline

Technology and Labor Demand: New Tasks Implication

Effect of **new tasks** on labor demand $=$ Productivity effect $+$ Reinstatement effect

- Expands the set of tasks which humans have a comparative advantage;
- Reinstatement effect:
  - Reinstates labor into new tasks;
  - **Always** increase labor share.

Thus, always increase labor demand.

The Effect on Labor Demand: Multiple Industries

Let's calculate economy-wide wage bill

$$\text{Wage bill} = \text{GDP} \times \sum_i \text{Labor share sector } i \times \text{Share of value added in sector } i$$

- We consider a multi-sector economy:

Composition effect:

- sectoral reallocations of labor (across sector);
- Note: displacement and reinstatement effect capture the change within sector.

Technology and Labor Demand: Multiple Industries

- Thus, effect on labor demand (redux):

Effect of **automation** in sector i on aggregate labor demand $=$ Productivity effect $+$ Displacement effect $+$ Composition effect

Effect of **new tasks** in sector i on aggregate labor demand $=$ Productivity effect $+$ Reinstatement effect $+$ Composition effect

Empirical Decomposition

$$\text{Change in aggregate wage bill} = \text{Productivity effect} + \text{Composition effect} + \text{Substitution effect} + \text{Change in task content}$$

- change in total wage bill is measured from data;
- Productivity effect: sum of the contribution from various sources of technology to value added $\Rightarrow$ changes in GDP per capita;
- Composition effect: reallocation of value added across sectors $\Rightarrow$ change in value added weighted by labor share;
- Substitution effect: substitution between capital and labor;

Change in Task Content

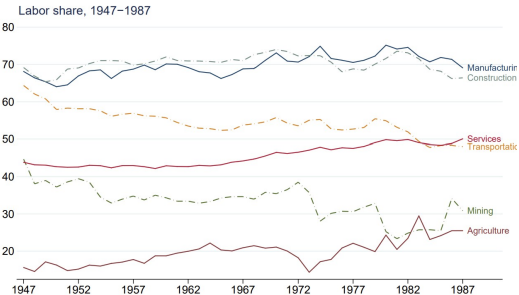
Under a competitive market,

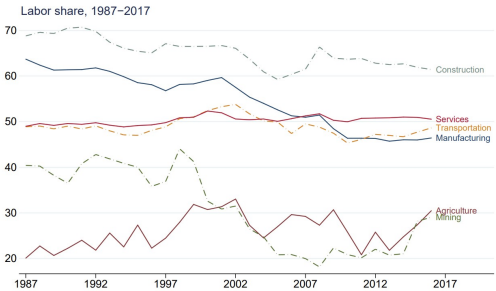
- change in labor share of one industry can be only explained by change in task content or substitution between labor and capital.
- whatever left unexplained by substitution effect is change in task content:

$$\begin{array}{c}\text{Change in task}\\ \text{content in } i\end{array} = \begin{array}{c}\text{Percent change in}\\ \text{labor share in } i\end{array} - \begin{array}{c}\text{Substitution}\\ \text{effect in } i\end{array}$$

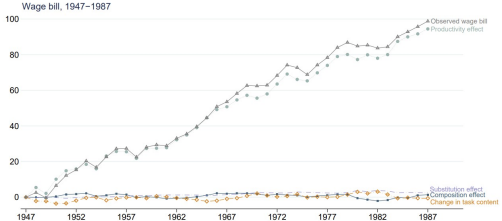
Thus, change in task content can be decomposed into **displacement effect** and **reinstatement effect**.

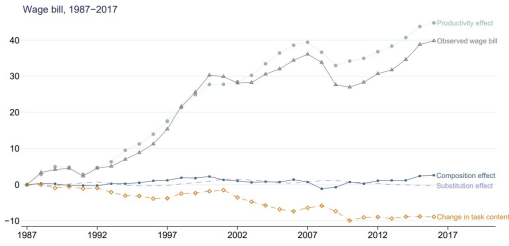
Evolution of Labor



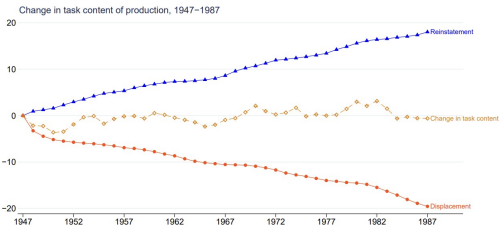


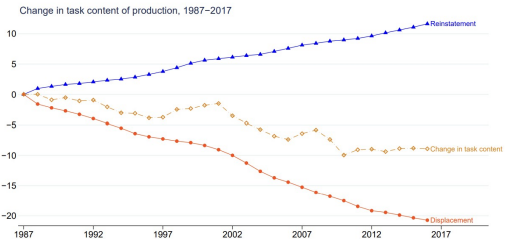
Wage Bill Decomposition





Change in Task Content





4. Early Impacts of AI

4.1 AI on Learning

The Cheating Incident

- Report
- X/Twitter thread
- The shooting incident at Brown on December 2025, had caused anxiety among students, and in Spring 2026 Prof. Serrano decided to give his students a take-home midterm for the first time in his career.
- The class is called Welfare Economics and Social Choice Theory, and in the past the average on midterm typically range from **65 to 80**.
- Because the promise of take home exam, this year this class attracts 86 students (up from 30 in the past years), and the average score on midterm this year is **96**.

The Cheating Incident Continued

- The professor suspected students used AI to cheat and he insisted switching to in person exam on the final.
- He announced that if the distribution of final grade is similar to the midterm, he will let it pass, otherwise, the midterm will be void, and the final accounts for 80% of the grade.
- Then **18** students dropped the class, and **9** didn't take the final, **3** people received 0, and the average on final is **48.6**, the lowest ever in his career.
- His previous final exam average never dipped below 65.

AI on Learning Penalty

- The important message from the story is not whether students' action can be accused of cheating;
- Rather, whether students delegated their learning process to AI, and eventually hurt students' capability to learn.
- As Prof. Serrano wrote: "We can't afford to become idiots".
- Not surprisingly, Stromberg et al. (2026) identified similar AI penalty on student learning.

The Generative AI Learning Penalty in Secondary School

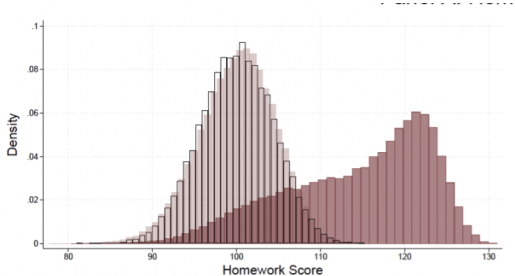
Stromberg et al. (2026):

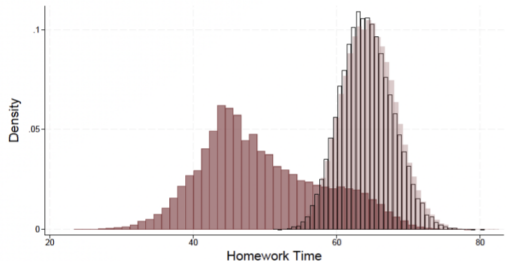
- Follow 26,811 7th to 12th grade students in central China to study **how self-directed use of AI affect their learning**.
- They follow students over 30 months, and link weekly homework scores and completion time to monthly closed book exams, county-wide exams, high school and college entrance examination scores.
- The data covers nine subjects.

Identification

- They track students exact date for AI tool registrations;
- They were able to know the precise timing of AI adoption.
- Students typically use AI tools such as: Doubao, Deepseek, etc.
- Because students adopt AI at different times, we can compare those **who adopted AI to the students who had not yet adopted** $\Rightarrow$ Difference in Difference.
- **The two groups had almost identical scores and follow the same trends**.
- This research design helps us to identify the causal effect of AI use on schooling learning.

Better Homework, Fast Completion



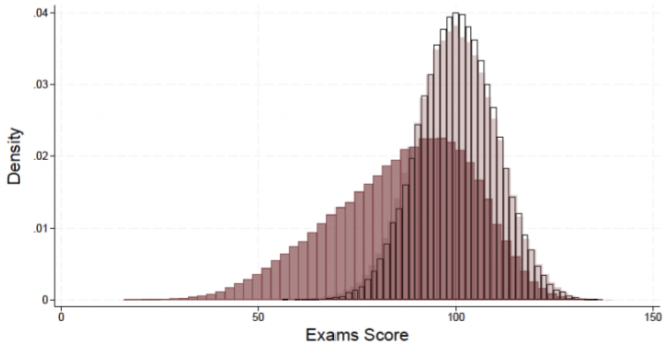


But Worse Exams

- AI indeed boost homework productivity: better homework score in less time spent on homework;
- But completing the homework is not the end goal for learning.
- We see a 20% decrease in closed book exams;
- And two years later, when it comes to high stakes exams like high school and college entrance examinations:
  - The penalty grow over time;
  - We see a 24% and 18% decrease in those exams, respectively.
- And the effect comes from the **compression of top performers**!

Exam Score Distribution

- Is AI adoption more on AI tutoring or homework outsourcing?



More on Learning

- So, maybe children and young adults are just more likely to be affected by AI adoption;
- They need time and guidance from parents and teachers to tell them how to use AI effectively, instead of outsourcing learning process to AI generated outputs.

Gross et al. (2026):

- U.S. STEM PhDs who use AI writing in dissertation are less likely to enter academia, and taking industry jobs instead.

Gross et al. (2026)

- The rapid advances of LLMs into science, raises concerns for delegating academic writing to AI **undermines** the academic training for future scientists.
- Thus, they study:
  - how heavily AI is used by PhD students;
  - who is using it;
  - how it relates to subsequent academic outcomes.

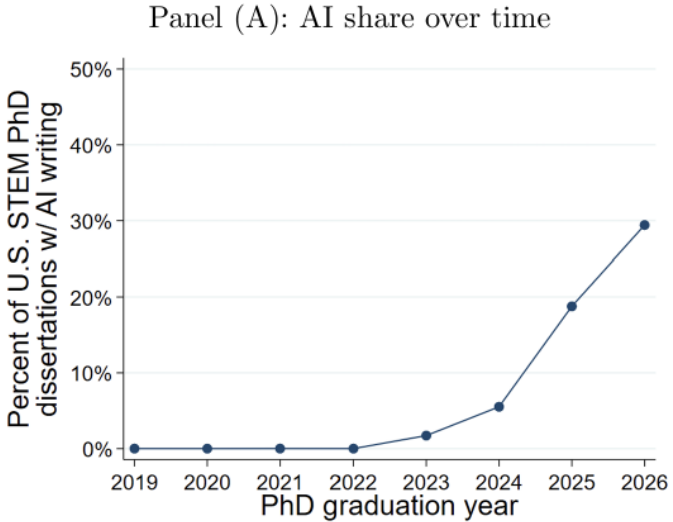
Data and Methods

- They collect near-all (226,000) STEM PhD students in the U.S. since 2019;
- Applied AI detection tools to dissertation abstracts;
- Linked individuals to LinkedIn profiles.

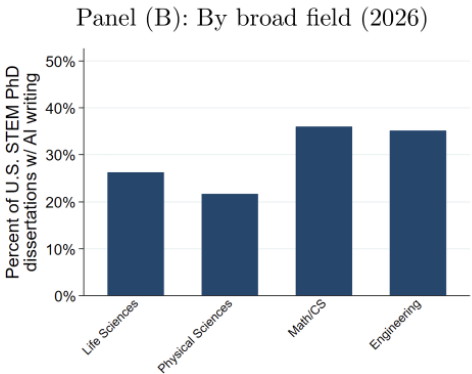
What about AI false detection?

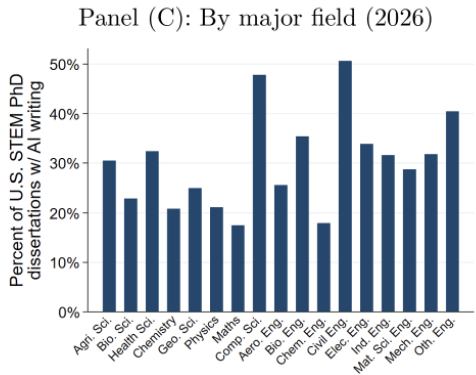
- Because prior to 2023, LLMs were not publicly available;
- They use pre-2023 data to validate the AI detection method.

AI Share over time



Computer Science and Engineering Field Used More

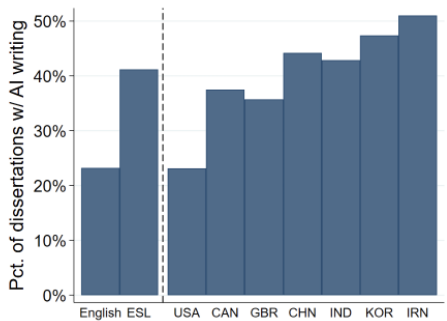




Country Origin

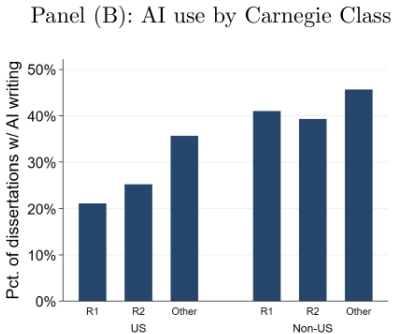
- They proxy country origin using the country of graduates' high school and bachelor degree.
- Non-native speakers used more AI in their dissertation.

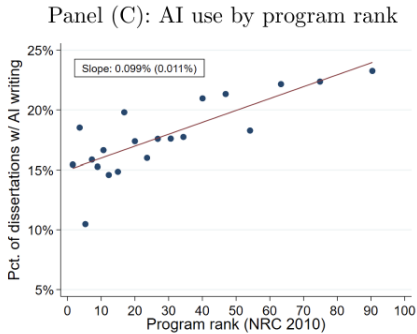
Panel (A): AI use by language and country



Program Ranking

- Dissertations remain human generated at the top.

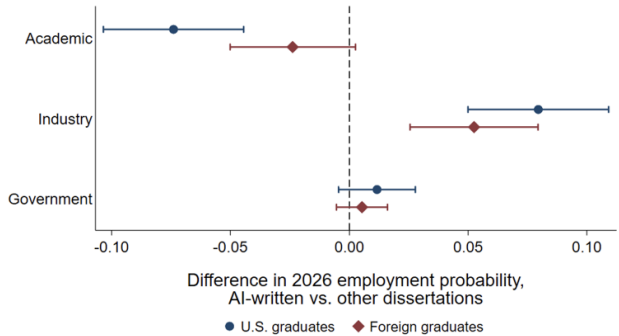




Graduating Out Of Academia

- AI writing appears to be common among those who exit academia.

Figure 3: Differences in employment patterns among graduates with AI-written dissertations



What about self selection?

- Students who decide to exit academia may use AI to expedite their research;
- But we see no effect on program duration.

More on selection bias:

- We still need to control students **academic potential**.
- For certain fields, industry jobs pay more on average;
- Positions like pharm postdoc may also have more research funding and opportunities;
- ...

Discussion

- The paper worried that whether delegating academic writing to AI itself hurts the essential training that help researchers succeed.
- Based on Acemoglu and Restrepo (2019), do you think that academic writing is one of tasks that will be replaced by AI in the future?

Discussion

- The paper worried that whether delegating academic writing to AI itself hurts the essential training that help researchers succeed.
- Based on Acemoglu and Restrepo (2019), do you think that academic writing is one of tasks that will be replaced by AI in the future?

The irony is:

- This paper itself was originally drafted by AI.

AI Draft Statement

We are grateful to Lee Fleming for his collaboration in related research, Kevin Bryan for helpful comments, and Ben Goihman for outstanding research assistance. We thank Pangram for API credits and discounted pricing, and the Duke University Fuqua School of Business, Alfred P. Sloan Foundation (Grant No. G-2023-21064), and National Science Foundation (Grant No. 2420824) for financial support. Claude Fable 5.1 was used in the initial drafting of the discussion of the empirical findings, and for a subset of the data appendix, though all AI-generated text has been substantially manually revised. All errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

4.2 AI on Labor Market

AI Exposure Index

- Anthropic AI
- Google AI Atlas
- OpenAI Signal

The Central Debate

As AI becomes more capable and widely used by firms, does it reduce or improve entry level hiring?

- Gates (2026) worry new grads will enter the workforce with fewer entry level jobs;
- Angelo (2026):
  - AI is just a scapegoat for downsizing;
  - AI until Dec 2025 was not good enough to replace the jobs they are cutting.
- Whether AI affects entry level jobs remain an empirical question.

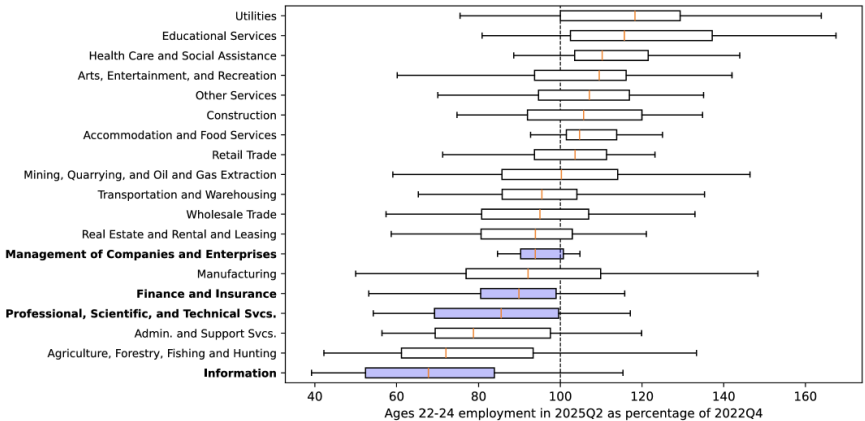
Tuker (2026) Show Differential Trends Among AI Exposed Jobs

Figure 3: Indexed employment by AI exposure quintile, ages 22-24



Impact by Sector

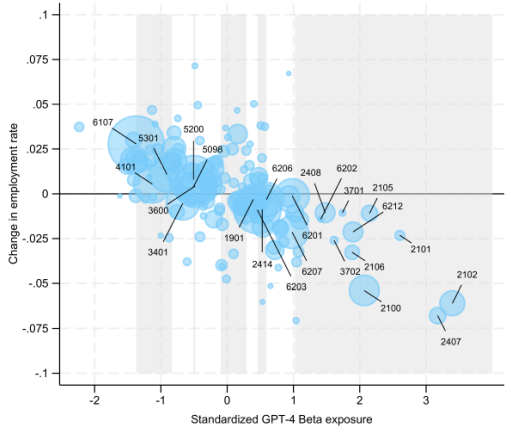
Figure 5: Distribution of changes in early career employment by sector



Orr et al. (2026) Further Explored the Impacts by College Majors

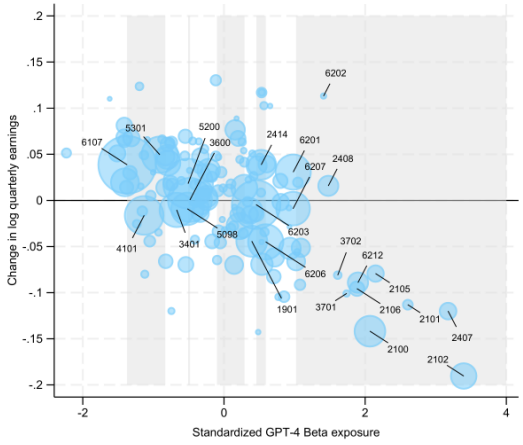
- They identify that the likelihood of initial employment and earnings dropped since the adoption of ChatGPT.
- The magnitude of initial earning loss is equivalent to **graduating into a large recession**.
- The effect attenuate as they move further from labor market entry, but remain large for most exposed majors.

Initial Employment



1901: Communications
2100: Computer and Info. Systems-General
2101: Computer Programming
2102: Computer Science
2105: Info. Sciences
2106: Computer Admin. Mgmt. and Security
2407: Computer Engineering
2408: Electrical Engineering
2414: Mechanical Engineering
3401: Liberal Arts
3600: Biology
3701: Applied Mathematics
3702: Statistics
4101: Physical Fitness, Parks, Rec., and Leisure
5098: Multi-disciplinary or General Science
5200: Psychology
5301: Criminal Justice and Fire Protection
6107: Nursing
6201: Accounting
6202: Actuarial Science
6203: Business Mgmt. and Admin.
6206: Marketing
6207: Finance
6212: Mgmt. Info. Systems and Statistics

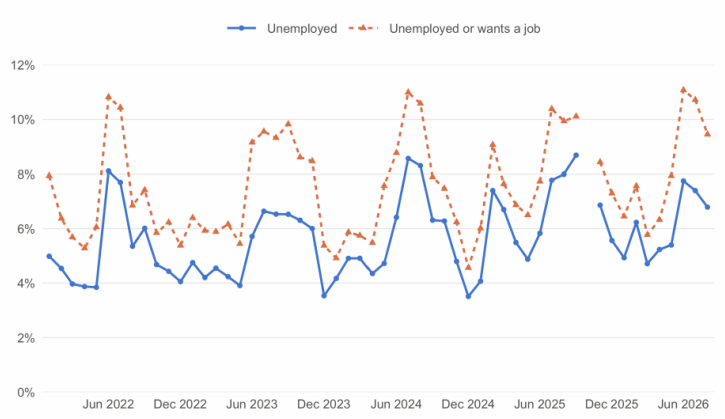
Initial Earnings



1901: Communications
2100: Computer and Info. Systems-General
2101: Computer Programming
2102: Computer Science
2105: Info. Sciences
2106: Computer Admin. Mgmt. and Security
2407: Computer Engineering
2408: Electrical Engineering
2414: Mechanical Engineering
3401: Liberal Arts
3600: Biology
3701: Applied Mathematics
3702: Statistics
4101: Physical Fitness, Parks, Rec., and Leisure
5098: Multi-disciplinary or General Science
5200: Psychology
5301: Criminal Justice and Fire Protection
6107: Nursing
6201: Accounting
6202: Actuarial Science
6203: Business Mgmt. and Admin.
6206: Marketing
6207: Finance
6212: Mgmt. Info. Systems and Statistics

Fairlie and Wu (2026) Found Nothing for 2026 Graduates

Figure 1. Unemployment among recent college graduates



Discussion

- Which task will become more rewarding in the future?
- The pandemic disruption makes the causal identification difficult empirically:
  - It could be the true impact of early AI impact;
  - Or downsizing the over expansion pre and during pandemic days.