

Jobs ‘at Risk From AI’ Were Already Declining Before ChatGPT Launched

Originally published at <https://www.unite.ai/jobs-at-risk-from-ai-were-already-declining-before-chatgpt-launched/>



A major new study finds jobs at risk from AI were already disappearing before ChatGPT launched, but students trained in those skills ended up with higher pay and quicker hires.

An extensive new research collaboration between US universities has found that the origins of the [AI-vulnerable jobs crisis](#) do not coincide with the launch of ChatGPT in late 2022, but rather that the issues begin earlier in the year, for apparently unrelated reasons.

Further, the report finds that graduates whose university majors were more ‘AI-exposed’ were actually associated with higher first-job salaries and shorter job searches after ChatGPT entered the market.

The new work leverages three large-scale datasets, including more than ten million scraped LinkedIn profiles, as well as unemployment records and insurance claims. The authors state:

‘[Our] results indicate that worsening labor-market outcomes in 2022–2024 for LLM-exposed workers and graduates were already underway prior to the mass-market emergence of LLM applications. Unemployment risk in highly exposed occupations rose beginning in early 2022—well before ChatGPT—and in most occupations and states we observe no discrete break coincident with its introduction.

‘Early-career workers were affected disproportionately: graduates from the 2021–2023 cohorts entered highly exposed jobs at lower rates and experienced longer observed delays to their first job than earlier cohorts, with gaps opening, again, before late 2022. At the same time, LLM-relevant education remained valuable within this environment.’

The new work re-frames the rise of AI as entering a job market *already weakened by broader economic and sector-specific pressures*, and notes that skills which complement AI kept, and may even have gained, market value.

The authors close out the paper by suggesting that ChatGPT’s November 2022 launch should not be treated as the terminator between the pre-AI and the AI-inclusive job market, and that a range of simultaneous circumstances should be considered along with the emergent influence of Large Language Models (LLMs):

‘These findings have implications for research and policy. First, they caution against treating ChatGPT’s launch as a clean natural experiment for AI’s labor-market impact: designs that attribute post-2022 labor-market weakness primarily to LLMs risk confounding AI diffusion with concurrent macroeconomic shifts (possible examples include monetary policy, sectoral demand, and/or post-pandemic adjustment).’

The authors suggest that universities and training programs should not drop skills often described as ‘vulnerable to AI’, such as writing, coding, and information synthesis. According to the results obtained in the work, teaching these skills in ways that work alongside AI, with an emphasis on checking outputs, judging quality, and using chatbots as tools rather than replacements, may help graduates to stay competitive in an unstable job market.

The [new study](#) is titled *AI-exposed jobs deteriorated before ChatGPT*, and comes from five researchers affiliated with various departments at the University of Pittsburgh, Stanford University, Chapman University, and Columbia University, in concert with Microsoft's AI Economy Institute at Redmond, and Revelio Labs in New York.

Method and Data


The paper's findings are, the authors note, a stark contrast to prior reports, including [one](#) from Stanford's Digital Economy Lab, as well as dark portents from luminaries such as the CEO of Anthropic, who [warned](#) in May of 2025 that AI 'could eliminate half of all entry-level white-collar jobs'.

The authors' analytical method initially examined unemployment among workers in occupations most exposed to AI-driven automation, with exposure defined using six-digit Standard Occupation Classification (SOC) codes, averaged to estimate exposure levels for broader two-digit SOC categories.

Monthly administrative data was drawn from the [ETA 203 report](#), compiled by the U.S. Department of Labor's Employment and Training Administration, detailing the most recent occupation of people claiming continuing unemployment insurance.

These data points were then combined with annual occupation-level employment figures from the Bureau of Labor Statistics' Occupational Employment and Wage Statistics program, allowing monthly unemployment risk to be estimated for each occupation within each state (where risk was defined as the likelihood that a worker in a given occupation claimed continuing unemployment benefits).

Historically, the paper notes, the jobs most exposed to AI faced 20-80% lower unemployment risk than less exposed roles, with the gap widening during the pandemic as remote-capable work proved more resilient. That advantage began to erode in early 2022, and by 2023–2024, the difference had largely disappeared:



Unemployment risk in AI-exposed jobs began rising in early 2022, ending a long period of relative stability. A shows this reversal as the gap between high- and low-exposure roles narrows before ChatGPT's launch. B reveals that the increase was concentrated in the most exposed quintile, with risk rising after a trough and then leveling off. C traces the effect to computer and math jobs, while most other fields remained stable. Risk was measured monthly across U.S. states and averaged quarterly. [Source](#)

As we can see in the graphs above, the authors grouped occupations into quintiles by 'AI exposure', and tracked these over time. Less exposed jobs consistently showed higher unemployment risk and stronger seasonal variation, with all groups peaking during the pandemic in 2020 and reaching a trough in early 2022.

After this low point, unemployment risk began rising in the most exposed quintiles, well before the launch of ChatGPT, and then stabilized, rather than accelerating in the months that followed.

Computer and math jobs saw the biggest rise in unemployment risk before ChatGPT launched, then leveled off. Most other roles showed little change. A few states, including California, Washington, and Alaska, saw post-ChatGPT increases, but national risk levels stayed close to pre-pandemic norms, indicating the influence of earlier economic pressures.


Data Considerations

The authors note that statistically, unemployment risk will reveal patterns across job types, but does not capture outcomes for specific groups – such as recent graduates who may not qualify for benefits or have a prior job to report. Other research and industry claims suggest that early-career workers face the greatest impact from AI, meaning overall unemployment data may miss those most affected.

To overcome this limitation, the new study drew on 10,584,980 LinkedIn profiles supplied by Revelio Labs. Each record from the dataset included detailed education histories covering degree type, field of study, graduation year, and university, alongside career data such as job titles (mapped to six-digit SOC codes), employers, start dates, and locations.

Job salaries were estimated using 'a proprietary machine learning model' trained on visa applications, self-reported entries, and public job postings, incorporating both role-specific details and individual career trajectories.

Since actual salaries could not be verified, the analysis also tracked the number of months graduates took to begin their first observed job within three years of finishing their degrees, excluding those with no recorded employment in that period (a metric which served as a proxy for labor market friction, assuming graduates update their profiles when hired):



Graduates entering the workforce after 2022 took longer to secure LLM-exposed jobs, but this decline in job market performance began months before ChatGPT's launch. Above, A shows that graduates with high-exposure first jobs typically found work faster, until this pattern reversed post-2022; B shows a similar delay for high-salary roles, though less pronounced; and C reveals that 2021 and 2022 cohorts entered LLM-exposed jobs at lower rates than earlier cohorts, with underperformance emerging before ChatGPT. Finally, D shows no equivalent shift for low-exposure jobs, reinforcing that the downturn predated widespread LLM adoption.

The authors analyzed job search duration across graduation cohorts, controlling for monthly job openings by state and sector, and accounting for differences in degree type, field of study, and university, with job exposure to LLMs defined using SOC codes.

Before ChatGPT's release, graduates entering highly exposed roles generally spent less time job seeking than their peers. For the 2023 and 2024 cohorts, this pattern reversed, with exposed roles taking longer to secure.

It should be emphasized that while the paper states that outcomes worsened after ChatGPT, the data show that this decline began months earlier and continued afterward, undermining the idea of a sudden post-ChatGPT collapse, and undermining also the attribution of the (ongoing) downward trend entirely to LLM uptake.

Educational Exposure

A central concern in the debate over AI and employment is whether students should continue training in skills that large language models may automate, such as writing, coding, or synthesis. If these skills have lost market value, then graduates most exposed to them should be faring worse. To test this, the authors estimated educational exposure to LLM-relevant tasks using LinkedIn profiles matched with millions of university syllabi, then tracked early job outcomes before and after ChatGPT:

Dependent Variable:	Log Salary (Real 2015 \$)		Months to First Job	
Variable	Model 1	Model 2	Model 3	Model 4
Job Opening Rate	0.028***	0.029***	0.473***	0.472***
Major SOC Log Annual Wage (2020)	0.726***	0.644***	-0.044***	-0.034***
Post ChatGPT Indicator (<i>gpt</i>)	0.034***	0.031***	2.341***	2.344***
Education Exposure (<i>edu</i>)	0.000	-0.000	0.000***	0.000***
<i>edu</i> × <i>gpt</i>	0.003*	0.001	-0.025***	-0.025***
LLM Exposure (<i>llm</i>)		0.564***		-0.096***
<i>gpt</i> × <i>llm</i>		-0.064***		0.034***
<i>edu</i> × <i>llm</i>		-0.019***		0.006***
<i>edu</i> × <i>llm</i> × <i>gpt</i>		0.012***		-0.009***
State FE	Yes	Yes	Yes	Yes
Sector FE	Yes	Yes	Yes	Yes
Terminal Degree Type FE	Yes	Yes	Yes	Yes
Graduation Year FE	Yes	Yes	Yes	Yes
College Major FE	Yes	Yes	Yes	Yes
University FE	Yes	Yes	Yes	Yes
R ²	0.426	0.433	0.374	0.374
adj. R ²	0.425	0.433	0.373	0.373
Number of Rows: 1,350,482				
P _{val} < 0.1*, P _{val} < 0.01**, P _{val} < 0.001***				

Educational exposure to LLM-relevant tasks predicts stronger early-career outcomes after ChatGPT. Post-2022 graduates with greater exposure to automatable skills were hired faster and earned higher salaries, partially offsetting the penalties linked to high LLM occupational exposure. All models control for job opening rates, job type, and educational background.

Before the advent of ChatGPT, this educational exposure showed no clear link to job search time or salary. After ChatGPT, it appears associated with faster hires and higher starting pay. Although roles with high LLM exposure tended to yield worse outcomes post-ChatGPT, graduates from more AI-aligned programs were less affected.

Rather than diminishing in value, skills seen as vulnerable to automation appeared to support better early-career outcomes.

'If LLMs were to blame for graduates' poor job market performance, then we would expect to see that education exposure indicates redundant skills that do not add value when job seeking.

'Yet, our results suggest that teaching AI-exposed skills yields better outcomes for graduates after ChatGPT's launch. These associations are difficult to reconcile with the view that LLM-relevant education became less valuable after ChatGPT. While not causal, they suggest that LLM-relevant preparation is at least compatible with better early-career outcomes in the post-ChatGPT period.'

The authors conclude by suggesting that the headline-garnering employment trends under study occurred in a labor market that was already being shaped by earlier events and trends. As it stands, separating the influence (on employment trends) of ChatGPT and AI in general from the unrelated forces that began the market downturn seems an impossible prospect, like trying to remove salt from soup.

** However, a fair amount of current commentary concedes that this kind of doomsaying from AI-invested founders is more akin to astroturfing, with the intent of dazzling potential clients and investors, and boosting stock prices.*

First published Wednesday, January 7, 2026

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.

Portfolio site: martinanderson.ai

Contact: martin@martinanderson.ai



Discover More