



Beyond the Binary: How Automation and Augmentation Are Combining to Reshape Work

Matt Sigelman, Benjamin Francis, Shrinidhi Rao and Gwynn Guilford

JANUARY 2026

Beyond the Binary: How Automation and Augmentation Are Combining to Reshape Work

Matt Sigelman, Benjamin Francis, Shrinidhi Rao and Gwynn Guilford

Suggested Citation: Sigelman, M. et al. (December 2025). Beyond the Binary: How Automation and Augmentation Are Combining to Reshape Work. Working Paper. Published by Burning Glass Institute.

Executive Summary

Is AI coming for your job—or making it better? Nearly three years into the ChatGPT era, we still don't have a clear answer. Headlines warn of mass displacement; techno-optimists promise a productivity renaissance. Meanwhile, the labor market sends mixed signals: entry-level hiring in AI-exposed fields has declined sharply, yet overall employment remains resilient. Workers report using AI to save hours each week, yet many fear their roles are next on the chopping block. The either/or framing—automation versus augmentation, job loss versus job creation—has dominated the conversation, but it may be obscuring what's actually happening.

Using labor market data from millions of job postings, we present early empirical evidence that cuts through this binary. What we find suggests a more nuanced reality: AI isn't simply replacing jobs or enhancing them. It's doing both—within the same roles, at the same time.

Our analysis reveals two key findings:

- **AI's predicted effects on skills are materializing in real time.** Automation-exposed skills were 16% more likely to see demand decline than baseline skills, while augmentation-exposed skills were 7% more likely to see demand increase. This divergence is highly statistically significant ($p < 0.0001$). These aren't dramatic swings—the labor market is adjusting, not cratering—but the pattern is clear and consistent across occupations.
- **Automation and augmentation are two sides of the same coin.** Conventional wisdom sorts jobs into those that will be automated and those that will be augmented. But the data show something different: the jobs experiencing the most automation are simultaneously experiencing the most augmentation. AI isn't eliminating project managers. It's transforming what project managers do each day.

These findings won't settle the debate over AI's long-term impact on employment—it's far too early for that. But they suggest the most urgent question may not be whether AI will take jobs, but how quickly the jobs it touches will change, and whether workers and institutions can adapt in time. That we can detect statistically significant effects this early—just a few years into widespread AI adoption—underscores the speed of the shift. It also highlights the value of real-time observation capabilities: systems that can track how skill demand is evolving as it happens, offering employers, educators, policymakers, and workers empirically grounded insights rather than speculation. The transformation is already underway; the open question is who will be ready for it.

I. Introduction: From Projection to Reality

When ChatGPT launched in late 2022, it triggered an avalanche of predictions about the future of work. Goldman Sachs **estimated** 300 million jobs could be affected globally. McKinsey **projected** that generative AI alone could increase the share of automatable work hours in the US economy by 8 percentage points by 2030. OpenAI's own researchers, in collaboration with academic economists, **concluded** that 80% of the US workforce could see at least 10% of their tasks affected by large language models.

Nearly three years later, we have something those projections lacked: actual labor market data.

This report presents early empirical findings on how AI is changing the demand for skills within occupations. Drawing on job posting data spanning five years—two years before ChatGPT's release and three years after—we measure shifts in employer demand for skills classified by their exposure to AI replacement and AI enhancement. The goal is not to forecast the future but to document what's already observable: the first measurable traces of AI's impact on what employers are looking for when they hire.

To be clear, this is not a story about AI being harmless. The disruption is real and already visible. Through October 2025, employers had cited AI as the reason for nearly 50,000 announced layoffs, making it the second-most common factor behind cost-cutting, according to **data** from Challenger, Gray & Christmas. Entry-level postings in AI-exposed occupations have dropped sharply, leaving young college graduates facing a much more challenging job market than older or less-educated peers, according to Burning Glass Institute **analysis**.

But displacement is not the whole story. Our analysis suggests that even as AI automates certain skills, it is simultaneously elevating demand for others—often within the very same roles. The result is not simply fewer jobs, but different jobs: roles that are being reshaped from within as the balance of tasks shifts.

This report proceeds as follows. Section II situates our work within the existing literature on AI and labor markets. Section III describes our data and methodology, including the LLM-based approach we used to classify skills by their AI exposure. Section IV presents our core findings. Section V discusses implications for workers, employers, and policymakers. Section VI concludes with reflections on what these early signals may—and may not—tell us about what lies ahead.

II. What We Know and What We're Adding

The question of whether AI will automate jobs or augment them is not new. Economists have debated automation's labor market effects for decades, and the emergence of large language models has only intensified the conversation.

Much of the early research on generative AI focused on exposure—identifying which occupations or tasks are most susceptible to AI capabilities. The influential “GPTs are GPTs” paper by Eloundou et al. (2023) mapped LLM capabilities onto occupational tasks, concluding that 80% of US workers could see at least 10% of their tasks affected. Felten et al. (2023) developed exposure scores linking language model capabilities to O*NET work activities. These frameworks provided valuable scaffolding for thinking about AI's potential reach, but they were projections, not measurements of actual labor market change.

More recent work has begun to distinguish between automation and augmentation effects. Handa et al. (2025) studied how millions of users actually interact with Claude.ai and discovered that the picture is more complex than simple substitution: in just over half of cases, workers engaged in back-and-forth dialogue with the AI to refine outputs, while in the remainder, they used it to complete tasks directly with minimal human involvement. The key finding was not the specific ratio, but rather that within individual occupations, workers deployed AI in both modes—sometimes as a thought partner, sometimes as a task executor.

Other researchers have taken a different approach. Brynjolfsson et al. (2025) studied the deployment of a generative AI tool among 5,172 customer service agents and found complex, heterogeneous effects within a single occupation. AI assistance increased overall productivity by 15%, but the impacts varied dramatically: lower-skill workers saw 30% productivity gains across multiple dimensions, while the highest-skill workers experienced small gains in speed but small declines in quality metrics. Rather than uniformly automating or augmenting the occupation, AI reshaped the distribution of performance within it.

Despite these advances, a gap remains in the literature. While Handa et al. document the mix of automation and augmentation in how AI is used, and Brynjolfsson et al. show heterogeneous impacts within an occupation, no study has yet systematically examined whether the same occupations experience both forces simultaneously when measured through employer demand for skills. The prevailing framework in public discourse—and implicitly in much research design—treats automation and augmentation as alternatives: jobs that get automated are distinct from jobs that get augmented.

Our contribution is to test that assumption against observed labor market data. Rather than asking whether individual tasks or interactions reflect automation or augmentation, we examine how employer demand for skills within occupations changes as AI capabilities expand. Using job posting data spanning five years, we measure changes in demand for skills classified by their exposure to AI replacement and AI enhancement, at the occupation level.

What we find challenges the binary framing. The jobs most exposed to automation are not a separate population from the jobs most exposed to augmentation. They are, in large part, the same jobs.

III. Data and Methodology

Our analysis draws on two core components: a large-scale dataset of job postings spanning the pre- and post-ChatGPT periods, and an LLM-based classification system that scores skills by their exposure to AI replacement and AI enhancement at the occupation level.

Job Posting Data

Job postings offer a real-time signal of employer demand that complements traditional labor market data. While employment statistics tell us who is working, postings reveal what employers are looking for when they hire—including the specific skills they request. This makes them a useful lens for detecting shifts in the composition of work, even before those shifts show up in aggregate employment figures.

For this part of the analysis, we sorted US job posting data into two periods within a five-year window: two years before ChatGPT's release (December 2020 through December 2022) and three years after (December 2022 through December 2025). This timeframe allows us to observe skill demand patterns before widespread LLM adoption and compare them to the period of rapid uptake that followed.

Classifying Skills by AI Exposure

To measure how AI is affecting skill demand, we first needed to classify skills by their relationship to AI capabilities. We developed an LLM-based approach that scores each skill, within each occupation, on two dimensions: its exposure to AI replacement (the degree to which AI can substitute for the skill) and its exposure to AI enhancement (the degree to which AI can amplify the skill's value or effectiveness).

The classification was conducted using a tournament-style methodology, in which an LLM evaluated skills pairwise to determine their relative exposure to replacement and enhancement. This approach yields continuous exposure scores rather than binary labels, allowing for more granular analysis.

This occupation-specific approach matters. The same skill may be highly automatable in one context and highly augmentable in another. "Writing," for instance, might be replacement-exposed for a role focused on producing routine documentation, but enhancement-exposed for a role where writing is a vehicle for strategic communication or persuasion. By scoring skills at the occupation level, we capture this contextual variation.

Measuring Demand Changes

With skills classified by their AI exposure, we measured how employer demand shifted between the pre-ChatGPT period (December 2020–December 2022) and the post-ChatGPT period (December 2022–December 2025).

Our approach compares high-exposure skills to a baseline of lower-exposure skills, separately for automation and augmentation. Specifically:

- **Automation analysis:** We identified skills in the top quartile of AI replacement exposure within each occupation. We then measured what share of these skills saw demand decline between the pre- and post-ChatGPT periods, compared to the share of skills in the bottom three quartiles that saw demand decline.
- **Augmentation analysis:** We identified skills in the top quartile of AI enhancement exposure within each occupation. We then measured what share of these skills saw demand increase between the pre- and post-ChatGPT periods, compared to the share of skills in the bottom three quartiles that saw demand increase.

Critically, we then examine whether these two forces operate within the same occupations. If the binary framing were correct—some jobs automated, others augmented—we would expect replacement-exposed and enhancement-exposed skills to cluster in different occupations. If instead they co-occur within the same occupations, this would suggest employers are simultaneously reducing demand for some skills while increasing demand for others in the same roles.

Limitations

Several limitations warrant acknowledgment. First, LLM-based classification involves judgment calls. While we designed the methodology to be systematic and replicable, the underlying scores reflect the model's interpretation of how AI capabilities map onto skills—an interpretation that may not perfectly match real-world dynamics.

Second, we are measuring early-stage effects. Three years of post-ChatGPT data captures initial signals, not long-term equilibrium. The patterns we observe may intensify, stabilize, or reverse as adoption matures and organizations adapt.

Third, job postings reflect employer demand, not actual job content. A skill's presence in a posting indicates that employers are asking for it—but not necessarily how that skill is used in practice, or whether AI is already augmenting or automating it behind the scenes.

Finally, correlation is not causation. While the timing of the shifts we observe aligns with LLM adoption, we cannot rule out other contributing factors. We present these findings as suggestive evidence, not definitive proof.

IV. Findings

A. AI's Effects on Skill Demand Are Already Measurable

The central question motivating this analysis is whether AI's predicted effects on skills are actually showing up in employer demand. The answer is yes—and the pattern is consistent with theoretical expectations.


We measured what share of skills saw demand decline (or increase) between the pre- and post-ChatGPT periods, then compared those shares for high-exposure skills versus baseline skills. The divergence is clear.

- For automation: 25% of automation-exposed skills saw demand decline in the post-ChatGPT period, compared to 21% of baseline skills—a gap of about 3.5 percentage points, or roughly 16% in relative terms.
- For augmentation: 61% of augmentation-exposed skills saw demand increase, compared to 57% of baseline skills—a gap of about 4 percentage points, or roughly 7% in relative terms.

This divergence is highly statistically significant ($p < 0.0001$), indicating that the observed pattern is extremely unlikely to be due to chance.

To be clear: we are not claiming that demand for automation-exposed skills has collapsed. Most of these skills did not see demand decline—but they were meaningfully more likely to decline than baseline skills. Conversely, augmentation-exposed skills were more likely to be growing. The effects are modest in magnitude but consistent in direction, and they align with what we would expect if AI is beginning to reshape the composition of work.

Figure 1. Demand Shifts for AI-Exposed Skills Relative to Baseline, Pre- vs. Post-LLMs



Source: Burning Glass Institute

The pattern holds across a range of occupations. Consider project management specialists: demand for skills like designing presentations and using Microsoft Excel—functions that AI tools can now handle efficiently—has softened relative to baseline, while demand for skills like budgeting and scheduling has strengthened. The role isn't disappearing; it's shifting. Similar patterns appear among sales representatives, where AI is absorbing operations and Microsoft Office-related tasks while amplifying demand for sales prospecting and customer relations management skills.

These are early signals, not definitive conclusions. However, the strength of these patterns—emerging just three years after ChatGPT's public release—suggests that AI's impact on skill demand is unfolding rapidly and in theoretically predictable directions.

B. Automation and Augmentation: Two Sides of the Same Coin


Finding A established that automation-exposed and augmentation-exposed skills are diverging in demand. But this raises a crucial question: are these two forces affecting different jobs, or the same ones?

Perhaps the most striking finding is not that automation and augmentation effects are both occurring—it's *where* they're occurring.

Conventional wisdom treats automation and augmentation as opposing forces that sort jobs into distinct categories. In this view, AI threatens routine-heavy roles—data entry clerks, telemarketers, basic bookkeepers—while enhancing high-skill professional work like software engineering, scientific research, or creative strategy. Policy discussions, workforce training programs, and career advice often reflect this binary: identify the at-risk jobs and steer workers toward the safe ones.

But the data tell a different story. When we calculate occupation-level AI exposure across both dimensions, we find that automation exposure and augmentation exposure are strongly positively correlated ($r = 0.87$, $p < 0.0001$, $n = 759$ occupations). The jobs most exposed to automation are also the jobs most exposed to augmentation. These are not separate buckets. They are, overwhelmingly, the same jobs.

Figure 2. Occupations Cluster Where Both AI Exposures Are High



Source: Burning Glass Institute

This makes sense in retrospect, even if it isn't immediately intuitive. Occupations with low AI exposure tend to be low on both dimensions—roles involving physical labor, manual trades, or highly unpredictable environments that AI cannot easily navigate. Occupations with high AI exposure tend to be knowledge work: roles rich in information processing, communication, and analysis, where AI can both substitute for routine tasks and amplify complex ones.

The implication is significant. AI isn't sorting the labor market into winners and losers. It's transforming job content across a broad swath of the economy. The project manager whose scheduling tasks are automated is the same project manager whose strategic responsibilities are expanding. The financial analyst who no longer builds models from scratch is the same analyst who now interprets and pressure-tests AI-generated outputs.

The relevant unit of analysis may not be the job, but the task. And the relevant question may not be "which jobs will survive?" but "how will this job change, and what new skills will workers need to perform it?"

V. Discussion

What These Findings Mean—and What They Don't

The patterns we observe are consistent with a labor market that is beginning to absorb AI's effects. Automation-exposed skills are softening; augmentation-exposed skills are strengthening; and the jobs experiencing both effects are largely the same jobs. This is not the mass displacement scenario that dominates headlines, nor is it the frictionless productivity boost that techno-optimists promise. It's something messier and more interesting: a gradual reshaping of work from within, where jobs persist but their substance transforms.

However, we want to be clear about the limits of these observations. Three years of post-ChatGPT data is a short window. The effects we measure are real, but they are early—and early signals can mislead. Adoption is still uneven across industries, firm sizes, and geographies. Many organizations are experimenting with AI without yet restructuring roles or adjusting hiring.

This uncertainty reinforces the value of real-time observation. The effects we document here may accelerate, plateau, or reverse; the only way to know is to keep watching. While periodic studies offer snapshots, what employers, workers, and policymakers need is something closer to continuous measurement.

We also cannot rule out alternative explanations. The post-ChatGPT period coincided with a tech hiring slowdown, elevated interest rates, and a broader recalibration of labor demand following the pandemic-era surge. Our methodology helps control for occupation-level trends, but it cannot fully isolate AI's causal effect. Longer-term data and more granular firm-level analysis will be necessary to disentangle AI's impact from broader economic shifts.

With those caveats in mind, we offer brief implications for three audiences.

For Workers: How to Stay Relevant

The finding here should offer some reassurance: so far, there isn't much empirical evidence for an AI firestorm. Jobs are not yet vanishing en masse. But they are changing, and the workers who thrive will be those who change with them.

The co-occurrence of automation and augmentation within the same roles means that staying relevant isn't about picking the right occupation or betting on the right skills. It's about paying attention to how your own job is evolving: which tasks are becoming easier, which are becoming more valuable, and where you should be deepening your expertise.

The practical implication is straightforward. Track the changes in your field. Acquire new skills where AI is creating demand. And double down on the capabilities that drive the greatest value in your role—particularly those that require judgment, context, or relationships, where AI remains a complement rather than a substitute.

For Employers: Redesigning Roles, Not Just Reducing Headcount

If the same roles are experiencing both automation and augmentation, the real opportunity lies in role redesign: rethinking not just headcount but job content. What should project managers actually spend their time on when AI handles scheduling and status updates? These are design questions, not just efficiency questions.

There's also a risk worth noting. As [prior Burning Glass Institute research](#) has shown, the decline in demand for automation-exposed skills is hitting entry-level roles hardest. Junior workers have traditionally built expertise by performing foundational tasks—drafting reports, building financial models, conducting initial research—the very tasks that AI now handles. If organizations aren't intentional about creating new pathways for skill development, they may find themselves with a hollowed-out talent pipeline and without the means of developing the very mid-career experts who will be increasingly necessary.

For Policymakers: Beyond the Displacement Frame

If AI is transforming jobs rather than simply eliminating them, then the policy challenge is not only cushioning job loss but facilitating job change. Traditional policy approaches have focused on identifying “safe” versus “at-risk” occupations—with an eye toward retraining workers to move into protected categories. But if automation and augmentation co-occur within the same jobs, then the challenge is not to reallocate workers across occupations, but rather to help existing workers adapt as their jobs transform.

This means investing in systems that support continuous adaptation mid-career: accessible retraining programs that focus on skill updating rather than wholesale career switches, portable credentials that recognize evolving competencies, and labor market information that helps workers understand how their specific roles are changing. The goal is not to move people out of AI-exposed occupations, but to help them succeed within those occupations as the work itself evolves.

It also means paying attention to who bears the adjustment costs. Younger workers in AI-exposed white-collar roles are facing disproportionate headwinds. The traditional logic—that education protects against technological disruption—may be inverting, at least temporarily. If college-educated knowledge workers require the same level of labor market support as previous waves of displaced manufacturing workers, our policy infrastructure is not ready.

VI. Conclusion

We set out to measure what has mostly been predicted: whether AI is changing the demand for skills, and how. The answer, at least in these early data, is yes—and in ways that complicate the prevailing narrative.

The labor market is not sorting neatly into automated jobs and augmented jobs. It is transforming jobs from within, automating some tasks while elevating others, often in the same roles. The project manager, the financial analyst, the sales representative: none of these jobs shows any signs of disappearing, but they are becoming different jobs. The skills employers are asking for are already shifting, and the shifts align with what we would expect if AI is indeed reshaping the composition of knowledge work.

This does not mean the disruption is painless. Entry-level workers in AI-exposed fields are already feeling the effects, and the adjustment costs are not evenly distributed. But the story is not simply one of replacement. It is one of transformation—messier, slower, and more varied than the headlines suggest.

What happens next is not predetermined. The trajectory depends on how quickly AI capabilities advance, how organizations choose to deploy them, and whether workers and institutions can adapt in time. We are still in the early chapters, and these early data are more useful for calibrating expectations than for making confident predictions.

But that calibration matters. The fact that we can detect statistically significant shifts this early—just three years into widespread AI adoption—suggests these changes are moving fast enough to warrant ongoing observation. Employers, educators, policymakers, and workers need more than periodic studies; they need real-time visibility into how skill demand is evolving so they can respond to what is actually happening rather than what forecasters predicted years ago. If the coming decade is less about jobs vanishing and more about jobs changing, then the questions we ask—and the investments we make—should reflect that reality.