SPECIAL COMPETITIVE STUDIES PROJECT | NVIDIA

A JOINT TASK FORCE PARTNERSHIP

AI AND THE FUTURE OF WORK

PRELIMINARY FINDINGS

Co Chairs

Senator Mike Rounds
U.S. Senator (R-SD)

Senator Mark Warner
U.S. Senator (D-VA)

Chris Malachowsky
NVIDIA

Ylli Bajraktari
SCSP

May 2026

AI and the Future of Work

Preliminary Findings

Letter from Task Force Co-Chairs

We stand at a genuine turning point for America's future. Artificial intelligence (AI) offers near unlimited potential to unlock new discoveries and tackle many of our most pressing problems, including in healthcare, agriculture, advanced manufacturing, and other areas. However, as AI reshapes many sectors of the American economy, it is raising major implications for the American workforce. Among other things, the country is undergoing changes to how knowledge is acquired, skills are developed, work is performed, careers are built, and opportunity remains available. It is critical that we understand these changes and respond with wise policy choices. The decisions made in the next few years about AI adoption, workforce investment, and education reform will shape the character of the American economy for a generation. We do not have the luxury of waiting for certainty before we act.

This Preliminary Findings report represents the Task Force's initial assessment of that challenge and of the opportunity that exists within it. This assessment is intended to inform a forthcoming final report that will propose policy recommendations for industry, academia, and officials at all levels of government. Our findings confirm that AI is not mere automation. It is a general-purpose technology with uniquely vast and still unpredictable impacts on labor, on mobility, and on the institutions Americans have long relied on to deliver economic security. Preliminary evidence gives both reasons to be hopeful and some signs of concern. On one hand, and after a long period of time, the U.S. economy is showing signs of an increase in productivity. On the other hand, AI-exposed occupations are under some pressure. Entry-level jobs, the traditional on-ramps to professional development and upward mobility, have seen some contraction in certain industries, although the extent of AI's direct role in the contraction remains debated. These are not abstract risks. They are gaps that, left unaddressed, could compound. Yet, the data infrastructure and strategies we would need to track and respond to these effects in real time does not yet exist at the necessary scale or quality.

However, the story is not one of inevitability. Our work has given us strong reason to believe that augmentation through AI that empowers workers, not displacement, can be the dominant paradigm of this transition, if the right choices are made and effective actions are put into place now by all. The window for proactive action is open. It will not remain open indefinitely.

That is why we convened this Task Force and why we are releasing these preliminary findings before our final report. Waiting for a perfect picture is itself a choice, and not the right one. The task before American firms, policymakers, educators, and workers is to meet the challenge and potential presented by AI, and to build the training pipelines, the data systems, the institutional partnerships, and the policy frameworks that allow this moment to expand opportunity rather than concentrate it.

We are grateful for the bipartisan support this effort has received, and for the engagement of the dozens of experts, practitioners, and leaders whose contributions have shaped this work.

The United States has met technological transitions before. It has done so not by managing decline or by surrendering to the inevitable, but by innovating and adapting, and embracing new opportunities that arise during those transitions. That is the task before us again and it begins with honest analysis, shared purpose, and a willingness to act before the moment passes.

Senator Mike Rounds Co-Chair

Senator Mark Warner Co-Chair

Chris Malachowsky Co-Chair

Ylli Bajraktari Co-Chair

About the Task Force

The Task Force on AI and the Future of Work was created to examine one of the most consequential missions of our time: ensuring that the United States not only leads the world in artificial intelligence, but that American workers are at the center of that leadership — equipped with knowledge, skills, opportunities, and tools needed to thrive in the age of AI.

The Task Force is a partnership between the Special Competitive Studies Project (SCSP) and NVIDIA. SCSP is a nonpartisan, not-for-profit project focused on strengthening America's long-term competitiveness as artificial intelligence (AI) and other emerging technologies are reshaping our national security, economy, and society. NVIDIA is the world leader in AI and accelerated computing. This partnership reflects a shared conviction that sustaining American leadership requires more than technological innovation. It requires a workforce ready to harness it and reap its benefits.

To pursue that mission, the Task Force recruited commissioners from across the full range of institutions that will need to act: U.S. Congress, universities, technological companies, philanthropies, and workforce development practitioners working on the front lines of economic transition. Its work is organized around three core objectives: defining the skills and educational frameworks needed for the jobs of the future; building a strategic roadmap that connects employers, educational institutions, and government to retrain workers for AI-augmented roles; and developing policies that sustain America's global leadership in AI innovation while ensuring its benefits reach broadly across the workforce.

Task Force Leadership

Senator Mike Rounds (R-SD), *Co-Chair*
Senator Mark Warner (D-VA), *Co-Chair*
Chris Malachowsky, Co-Founder of NVIDIA, *Co-Chair*
Ylli Bajraktari, President of SCSP, *Co-Chair*

Task Force Members

Dr. Erik Brynjolfsson, Stanford University
Dr. France Cordova, President Science Philanthropy Alliance, External Faculty Professor, Santa Fe Institute
Dr. José-Marie Griffiths, President Dakota State University
Eric Holcomb, Former Governor of Indiana
Dr. Tom Mitchell, Carnegie Mellon University
Chan Park, Head of U.S. and Canada Policy and Partnerships, OpenAI
Gina Raimondo, Former Secretary of Commerce

Preliminary Findings

These preliminary findings provide an initial assessment by the Task Force of how AI is beginning to reshape work. They are preliminary because the data and the evidence are still emerging, and there is still a great degree of directional uncertainty. Nevertheless, these findings offer an initial diagnosis of AI's impact on the character of work, identify some early signals, and map out both opportunities and risks emerging for our economy and American workers. Together, they are intended to ground the Task Force's subsequent public policy recommendations in a coherent and empirically informed framework. The Task Force aims to arrive at a final report with such recommendations later this year.

1. AI is Advancing at Unprecedented Speed and More Broadly Than Prior Technologies

AI is a general-purpose technology with an unusually rapid cycle of improvement and adoption across sectors, diffusing across the United States twice as fast as the internet. Unlike most prior general-purpose technologies, it also extends into certain types of cognitive work, reshaping not only how some tasks are performed but the foundational skills with which those tasks are performed. This combination of speed and scope increases both the scale of potential gains and the uncertainty around the potential near-term disruption.

2. Historically, Broad-based Technological Change Has Reshaped Work, but the Exact Direction and Full Extent of AI's Impact is Difficult to Predict

History shows that technological revolutions simultaneously create, transform, and eliminate jobs—but with outcomes that are notoriously difficult to predict in advance. While past general-purpose technology advances enhanced productivity, created new industries, and ultimately expanded work, they also involved unpredictable and uneven adjustment periods and resulted in fundamental, permanent shifts in many occupations. As with past precedent, AI has the potential to expand the frontier of human work and result in net positive job creation, but uncertainties during the transition period are likely to create disruption for at least some workers and complicate projections of the near-term impacts.

3. Measuring AI's Impact on Labor Market Requires a Fundamentally Different Data Collection Strategy

Existing labor market data is not well suited to capture the speed and the breadth of changes that AI is already driving. Key indicators such as employment and productivity are generally lagging indicators that do not capture the underlying shifts in work in real time. As a result, early effects—such as changing skill demands or hiring patterns—are difficult to observe with traditional metrics and will require new data collection and analytic methodologies.

4. The Impact of AI is Concentrated at the Task Level First, not at the Occupation Level

Jobs are composed of bundles of tasks, and where AI is being adopted its impact appears to be concentrated in automating or augmenting the performance of specific tasks rather than eliminating entire occupations. This leads to a gradual but potentially fundamental reconfiguration of roles, often without immediate changes in job titles or employment levels. Therefore, it is increasingly vital for workers to understand how AI impacts key tasks in their roles, the complementary tasks that gain prominence, and the new necessary skills to enhance their individual productivity.

5. Preliminary Labor Market Signals Indicate AI's Pressure on Entry-Level Roles

Aggregate employment data remains generally strong; however, preliminary evidence suggests a contraction in entry-level roles in AI-exposed occupations over the last several years. While the extent of AI's direct role in the downturn remains debated, these positions may be particularly vulnerable to impact from AI because they involve routine and largely standardized tasks that AI may be able to perform. Whether this long-term impact on entry-level jobs is a bellwether to a broader impact to come is currently hard to ascertain. Nevertheless, the implications remain significant, in part because AI may disrupt the pipeline of talent and the well-established pathways for career progression in certain industries.

6. The Extent of AI Adoption Will Be One of the Principal Determinants of the Impact on Labor Market Outcomes

The potential impact of AI on the economy appears to depend on what the technology can do. However, one must not overlook the importance of adoption, which is neither pre-ordained nor even. Adoption varies across sectors and organizations, shaped by incentives, workflows, and institutional constraints. The balance between augmentation and automation will be determined by these implementation choices, and who participates in making them, rather than technological inevitability.

7. Shifting Skill Demand in the AI-Driven Economy

As AI has started to reshape the skills required in the economy, it is inevitably raising questions about the type of training and roles that may have more enduring value. In some AI-exposed occupations, new digital and technical skills, as well as durable human capabilities—such as critical thinking, leadership, and adaptability—may be key to prosperity. In other emerging and expanding industries—such as new AI-focused job categories and construction to support AI infrastructure—the supply of skilled labor is insufficient.

8. Traditional Educational Pathways and Trainings May Fall Short in Meeting the AI Moment

The current education, credentialing, and training systems remain largely built around front-loaded learning in the lead up to a degree and career, leaving them misaligned with rapidly changing skill demands. To keep pace, these systems must be reimagined to leverage AI as a training mechanism and to produce a balanced mix of foundational capabilities and adaptable domain-specific skills.

9. The Distribution of AI's Gains Need Not Be Concentrated, and Can Be Shaped by Implementation Choices

AI can broaden access to high-productivity tools and raise worker performance, but it can also concentrate gains among firms, regions, and workers with early access and complementary skills. The outcome will depend on how AI is deployed, who benefits from productivity gains, and how quickly workers can transition skill sets and jobs.

10. The Need for Coordinated, Whole-of-Nation Action

No type of institution—government, industry, or educational—can manage the AI-powered transition alone. The speed, scale, and complexity of AI's impact require coordinated action across sectors and actors to align incentives, invest in workforce development, and ensure broadly shared benefits. The window for proactive intervention is open.

1. AI is Advancing at Unprecedented Speed and More Broadly Than Prior Technologies

Artificial intelligence is widely regarded as a transformational technology poised to reshape the workforce in ways comparable to earlier general-purpose innovations such as electricity, the railroad, and the internet.[1] Like those technologies, AI has the potential to diffuse across nearly every sector of the economy, fundamentally altering how work is organized, how productivity is achieved, and which skills are most valuable. The historical record provides clear examples of this profound effect. Electricity enabled factories to reorganize production and extend operating hours,[2] railroads expanded markets and integrated regional economies into national supply chains, and the internet digitized communication, lowered transaction costs, and accelerated the flow of information. AI similarly promises to reconfigure workflows, automate certain tasks, and enable entirely new categories of products and services. As with past technological revolutions, the transition is likely to be uneven and potentially disruptive in the short term, but ultimately augmentative and systemic in its long-term effects on economic growth and labor demand.[3]

However, AI may have an even greater impact than earlier technologies because it does not merely augment or replace physical labor—it extends into the realm of knowledge work itself. Previous waves of automation primarily affected routine manual tasks, but AI systems are increasingly capable of performing activities that require reasoning, pattern recognition, language understanding, and decision-making. This means that jobs once considered insulated from automation—such as those in law, finance, medicine, education, and creative industries—are now being reshaped. While there will be variance across sectors, it may not be that AI is simply replacing workers, but reconfiguring roles by handling specific cognitive tasks, thereby changing how professionals allocate their time and expertise. This shift blurs the traditional boundary between "automatable" and "non-automatable" work, making AI a uniquely far-reaching force in the labor market.

The consequential scope of AI is amplified by its speed. Compared to the personal computer or the internet, generative AI has seen faster mass adoption, with 53% of individual respondents and 88% of organizations reporting using AI in a global survey in 2025, less than 3 years after the public launch of ChatGPT—twice as fast as the internet.[4] This use is continuing to spread and

---

[1] Tyna Eloundou, et al., GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models, Arxiv (2023).
[2] Paul David, The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox, American Economic Review (1990).
[3] Joseph Briggs & Devesh Kodnani, The Potentially Large Effects of Artificial Intelligence on Economic Growth, Goldman Sachs (2023).
[4] Artificial Intelligence Index Report, Stanford University Human-Centered Artificial Intelligence (2026).

deepen as the capability of AI models is increasing rapidly, leaving the tests created to measure their progress obsolete. The increasing performance and adoption of generative AI is reflected in the consumer surplus these tools provide, which hit $172 billion in 2026, a 54% increase from the previous year.[5]

A growing body of research underscores the scale of AI's potential impact. Studies from organizations such as McKinsey,[6] the IMF,[7] and academic institutions[8] estimate that a significant share of current work tasks—between 20% and 50%—could be automated using existing or near-term AI technologies. Research also highlights the dual nature of AI's impact: while it may automate some existing tasks within jobs, it is expected to increase the demand for other tasks where humans retain a comparative advantage.[9] The impact of AI will be lower in the near-term in roles that rely heavily on capabilities that require contextual judgment,[10] social intelligence,[11] ethical reasoning, and physical dexterity—such as service industries, skilled trades, and high-judgement roles. At the same time, the deployment of AI will create new products and services and generate demand for entirely new tasks that were previously infeasible, thereby expanding the frontier of human work. Ultimately, the effect on employment may depend heavily on how quickly institutions—such as education systems, labor markets, and firms—adapt to these changes.

Given these dynamics, fully harnessing the potential of AI is not merely an economic opportunity but a strategic imperative for the United States. Leadership in AI adoption and innovation will play a central role in determining global competitiveness, influencing everything from industrial productivity to national security. Countries that effectively integrate AI into their workforce—through investments in education, reskilling, research, and infrastructure—will be better positioned to drive scientific discovery, technological advancement, and economic growth. Conversely, failure to adopt AI could result in slower productivity gains and erosion of competitive advantage. Ensuring that workers can successfully transition alongside AI, while fostering an environment that supports innovation, will be essential to maintaining America's leadership in the global economy.

---

[5] Artificial Intelligence Index Report, Stanford University Human-Centered Artificial Intelligence (2026).
[6] Joseph Briggs & Devesh Kodnani, The Potentially Large Effects of Artificial Intelligence on Economic Growth, Goldman Sachs (2023).
[7] Mauro Cazzaniga, et al., Gen-AI: Artificial Intelligence and the Future of Work, International Monetary Fund (2024).
[8] Tyna Eloundou, et al., GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models, Arxiv (2023).
[9] David Autor, Why are there still so many jobs? The History and Future of Workplace Automation (2015).
[10] David Autor, Polanyi's Paradox and the Shape of Employment Growth (2014).
[11] David Rizzo, et al., Human Capabilities are at the Heart of High-Performing Teams (2026).

2. Historically, Broad-based Technological Change Has Reshaped Work, but the Exact Direction and Full Extent of AI's Impact is Difficult to Predict

Historical Discussion of Technological Movements

Historical patterns of technological change and labor markets reveal three key insights for the future of work in the age of AI. First, technological change has a varied impact on jobs and the tasks that compose those jobs, automating some tasks while enhancing others through augmentation. This process is often described as 'Creative Destruction'—the essential fact of capitalism where the innovative entry of new technologies incessantly revolutionizes the economic structure from within, morphing the old structure into a new, more efficient one.[12] Second, technological change has historically also created entirely new categories of jobs that could not have existed prior to the invention of the new technology. Third, technological impact on the workforce often takes longer than predicted because absorption and reconfiguration take time; jobs will change and new ones will be created, but it is difficult to identify ahead of time what those might be. The total number of jobs in an economy is not fixed; the supply of jobs is not a finite pool to be divided among workers, but can expand when technology and policy create the conditions for new industries and forms of employment to emerge.[13] Technically, 60% of work in the United States today is "new work" since 1940 — types of employment that did not exist before that date.[14] These jobs are a result of technological advances, like software engineers, but also incorporate previously uncategorized or unestablished jobs, including conference planners or nail technicians.[15]

Historical Examples

A key historical example is the first Industrial Revolution (1760-1840). This period is one of the most commonly drawn parallels to AI today. Industrialization in the West accelerated production through mechanization, concentrated wealth among capital owners, and upended worker norms. The working and agrarian classes felt left behind by automation and disenfranchised in their ability to earn wages, as seen with the Luddites who destroyed the machines replacing them in protest.[16] Most weavers, for example, lost their jobs due to automation or consolidated into singular roles managing new machines. But new,

---

[12] Ramazan Uctu, et al., Creative destruction and artificial intelligence: The transformation of industries during the sixth wave, Journal of Open Innovation: Technology, Market, and Complexity (2024).
[13] Lawrence Marsh, Automation and the Lump of Labor fallacy, Notre Dame Economist (2025).
[14] David Autor, et al., New Frontiers: The Origins and Content of New Work, 1940–2018, The Quarterly Journal of Economics (2024).
[15] Peter Dizikes, Most Work Is New Work, Long-Term Study of U.S. Census Data Shows, MIT News (2024).
[16] Daron Acemoglu & Simon Johnson, Learning from Ricardo and Thompson: Machinery and Labor in the Early Industrial Revolution, and in the Age of AI, NBER Working Paper 32416 (2024).

complementary jobs were created such as machine operator and mechanic. While there was a downtick in roles in the short to medium term, as well as a societal shift, the benefits of industrialization are undeniable.

The age of electricity (1870-1914) brought another significant wave of change with the invention of the lightbulb, in-house adoption of electricity, necessity of power grids, and large-scale transportation. While the lamp lighter job was automated with the influx of electric lamp posts, jobs like train conductor were augmented. Growth within this period for certain jobs was exponential: in 1870, there were 55,000 machinists and in 1900 there were 283,000.[17] Jobs like machinists and electricians were multiplied because they were complementary to the original invention. With strong general-purpose technologies, an ecosystem emerges around them that creates opportunity. Most, if not all, jobs today are dependent on electricity and it has significantly improved quality of life and work.

Even more recently, Americans have been wary of technological eras and have turned to policy for support. In 1964, President Lyndon Johnson signed into law the "National Commission on Technology, Automation, and Economic Progress."[18] This bill addressed concerns about the pace of technological change, its results on employment, and the effect on society and individual lives. Similarly to today, Americans feared job replacement and that society could not keep up with the rate of change, leaving Americans without jobs and purpose. President Johnson advocated that "automation can be the ally of our prosperity if we will just look ahead, if we will understand what is to come, and if we will set our course wisely after proper planning for the future."[19] An optimistic attitude towards the future and structured approaches to grow alongside the technology have proved effective through the 1960s and in prior historical case studies.

The Historical Record: Disruption and Resilience

The historical record is, on balance, a story of resilience. Specific occupations disappear—weavers displaced by the power loom, lamplighters rendered obsolete by electrification—but the broader arc of technological change has consistently produced more work, not less. Each major transition was met with the same fear: that this disruption was too vast and too fast for society to absorb. And each time, that fear proved wrong. New categories of work emerged that could not have been anticipated in advance, and the economies that embraced change ultimately produced higher wages, broader prosperity, and better working conditions than

[17] David Montgomery, Chapter 3: Labor in the Industrial Era, U.S. Department of Labor (last accessed 2026).
[18] Public Law No. 88-444, An Act to establish a National Commission on Technology, Automation, and Economic Progress, (1964).
[19] Lyndon B. Johnson, et al., Remarks Upon Signing Bill Creating the National Commission on Technology, Automation, and Economic Progress, The American Presidency Project, (1964).

those that preceded them. We may already be witnessing this through the emergence of new AI job categories and the escalating demand for labor to support infrastructure needs, topics further explored in section seven.

The pattern across the Industrial Revolution, the age of electricity, and the computing era is remarkably consistent. General-purpose technologies create ecosystems. The steam engine required mechanics. Electrification required electricians and engineers. Computing required programmers, systems administrators, and an entirely new class of knowledge workers that did not exist a generation earlier. In each case, the jobs created by the technology far outnumbered the jobs it displaced, and the new roles were, on the whole, safer, better compensated, and more cognitively engaging than what came before. This does not suggest that the changes were smooth or that the pace of change was even; managing the displacement that inevitably will occur alongside job creation will require policy decisions based on lessons from periods of past technological change. Society must address economic displacement, anxiety, and the need for transition. History does not promise that every worker will land well in every transition. But it does offer a clear and repeated demonstration that technological change, over time, expands the frontier of human work rather than contracting it.

3. Measuring AI's Impact on the Labor Market Requires a Fundamentally Different Data Collection Strategy

Managing the transition to an AI-integrated economy requires better data—and better metrics to track its effects as they unfold. Conventional economic indicators were designed for a slower, more tangible economy. They are largely backward-looking and, as a result, poorly equipped to capture the pace, mechanisms, and distributional consequences of AI-driven change.

The Conceptual Gap

The core challenge is a fundamental disconnect between what AI changes and what standard metrics track. GDP measures market transactions, but AI delivers much of its value outside of them through free tools, subscriptions, and productivity gains that do not yet register in output figures.[20] Labor productivity tracks output per worker and should, in principle, capture AI-driven improvements, but as long as AI's contribution cannot be cleanly separated from other inputs, observed gains remain ambiguous.[21]

Labor market indicators face similar limitations. The headline U-3 unemployment rate tracks those without jobs who actively sought work in the prior month; at 4.4%, it reads as healthy.[22] The broader U-6 measure, which captures discouraged workers and those involuntarily working part-time, must also be considered as it often reacts quicker during periods of significant disruption. Existing labor market data systems are too fragmented and slow to track skills-based underemployment, and unable to link education, employment, and occupational outcomes in real time.[23]

To fill these gaps, researchers have developed two broad AI-specific measures that link model capabilities to labor market data. Exposure-based measures map AI capabilities onto occupational tasks to estimate which jobs are theoretically affected, capturing potential impact, not actual use.[24] Usage-based measures track real-world AI activity, offering a more direct window into how AI is being used, but they are limited to specific tools and user populations and

---

[20] Erik Brynjolfsson & Avinash Collis, How Should We Measure the Digital Economy?, Brookings Institution (2020).

[21] Erik Brynjolfsson, et al., GDP-B: Accounting for the Value of New and Free Goods in the Digital Economy, National Bureau of Economic Research Working Paper 25695 (2019).

[22] Board of Governors of the Federal Reserve System, Unemployment Rate, FRED, Federal Reserve Bank of St. Louis (2026).

[23] Briefing to the Task Force (March 27, 2026).

[24] The most widely used AI exposure indices each map AI capabilities onto occupational task descriptions to produce exposure scores by job type: Edward Felten, et al., Occupational Heterogeneity in Exposure to Generative AI, SSRN (2021); Tyna Eloundou, et al., GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models, arXiv (2023).

are therefore not fully representative.[25] The gap between what AI can theoretically do and the ways workers are actually using it remains a central open question for labor market analysis, and is examined in greater depth in Section 8.

The Temporal Gap

Aggregate data is retrospective by design; it captures what has already settled, not what is actively shifting. This lag is an inherent feature of national accounting and labor surveys and a familiar challenge in measuring major technological transitions. In 1987 Robert Solow observed that "you can see the computer age everywhere but in the productivity statistics," a paradox that took another decade to resolve as firms reorganized around computing.[26]

Economists describe the underlying dynamic as the productivity J-curve: initial technology investments lead to flat or declining measured output before gains materialize, as organizations build the complementary intangibles, new processes, skills, and business models needed to realize the technology's potential.[27] Given the recency and speed of widespread AI adoption, current macroeconomic indicators should be interpreted as lagging signals, not real-time measures of impact.

What the Data Shows, and What It Blurs

At the aggregate level, total United States employment has grown approximately 2.5% since the public launch of ChatGPT in late 2022, and unemployment remains low. Sectoral data shows that employment in computer systems design—one of the most AI-exposed industries—has declined by roughly 5% since late 2022, while the top decile of AI-exposed industries has seen modest employment contraction even as the broader economy expanded.[28] However, other data suggest renewed demand for some software development and computer-related roles, and BLS projects faster-than-average growth over the next decade for several of these occupations..[29] At any rate, these patterns are descriptive, not causal: the industries most affected by the Federal Reserve's rate cycle and the post-pandemic correction in tech hiring are also the most AI-exposed. Disentangling these forces is very difficult with available data.

The broader research literature reflects this difficulty: studies using different data sources, methodologies, and exposure measures have reached very different conclusions. This divergence is unsurprising as these approaches capture effects at different levels of analysis.

---

[25] The Anthropic Economic Index, Anthropic (2025).
[26] Robert M. Solow, We'd Better Watch Out, New York Times Book Review (1987).
[27] Erik Brynjolfsson, et al., The Productivity J-Curve: How Intangibles Complement General Purpose Technologies, American Economic Journal: Macroeconomics at 333–372 (2021).
[28] J. Scott Davis, AI Is Simultaneously Aiding and Replacing Workers, Wage Data Suggest, Federal Reserve Bank of Dallas (2026).
[29] Bureau of Labor Statistics, Occupational Outlook Handbook: Computer and Information Technology Occupations, U.S. Department of Labor (2026); .

Crucially, a technology that disrupts tasks within jobs before causing displacement will leave its earliest fingerprints in the changing composition of skills demanded by employers, not in headline employment figures. As a result, the most visible early effects—shifts in skill demands and entry-level opportunities—are only detectable in granular, micro-level data.

Emerging evidence from job postings, employer surveys, and internal hiring records is beginning to capture these shifts, and is examined in detail in Section 5. These indicators considered together—rather than single headline figures—provide the most reliable picture of where the transition currently stands. Some research ecosystems have begun addressing these gaps through data-sharing consortia that combine job postings, firm-level, and administrative data. Expanding such approaches could improve visibility into early labor market shifts.

4. The Impact of AI is Concentrated at the Task Level First, not at the Occupation Level

An occupation is a bundle or collection of discrete tasks performed by a single worker. A financial analyst's job, for instance, bundles data retrieval, model construction, interpretation, client communication, and judgment calls into a single role. Tasks are the more granular unit: discrete work activities whose execution can, in principle, be evaluated independently of who or what performs them. Technology substitutes for tasks first; only rarely later—and only sometimes—does substitution accumulate to the point of eliminating an occupation. Automation concentrates on specific tasks within a job where machine execution is more cost-effective than human labor. The result is generally a reconfigured occupation: the job title may remain the same, but the underlying task bundle is fundamentally altered. The lag inherent in this process is consequential; it means that the transformation already underway in white-collar knowledge work is largely invisible in aggregate employment statistics, which track headcounts by occupation rather than the shifting composition of tasks within those occupations. The effects of AI will show up in job titles only after they have already reshaped what the work involves. Estimates suggest that around 80% of the U.S. workforce could see at least 10% of their tasks affected, with exposure rising in higher-income occupations on average.[30]

Mechanisms of Change: Augmentation and Automation

AI reshapes the internal mix of tasks through two primary channels. Augmentation occurs when AI enhances worker productivity without replacing the worker—accelerating task completion, improving output quality and enabling less experienced workers to perform at levels previously associated with greater experience. Automation, by contrast, involves AI substituting for human labor in specific tasks, with two distinct economic outcomes. Displacement occurs when automation reduces the human hours required for a given activity; reinstatement occurs when the productivity and wealth gains from automation generate new complementary tasks where human labor retains comparative advantage—one of the mechanisms through which prior waves of technological change ultimately expanded rather than contracted total employment.[31]

---

[30] Tyna Eloundou, et al., GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models, arXiv (2023).

[31] Daron Acemoglu & Pascual Restrepo, Automation and New Tasks: How Technology Displaces and Reinstates Labor, Journal of Economic Perspectives at 3–30 (2019).

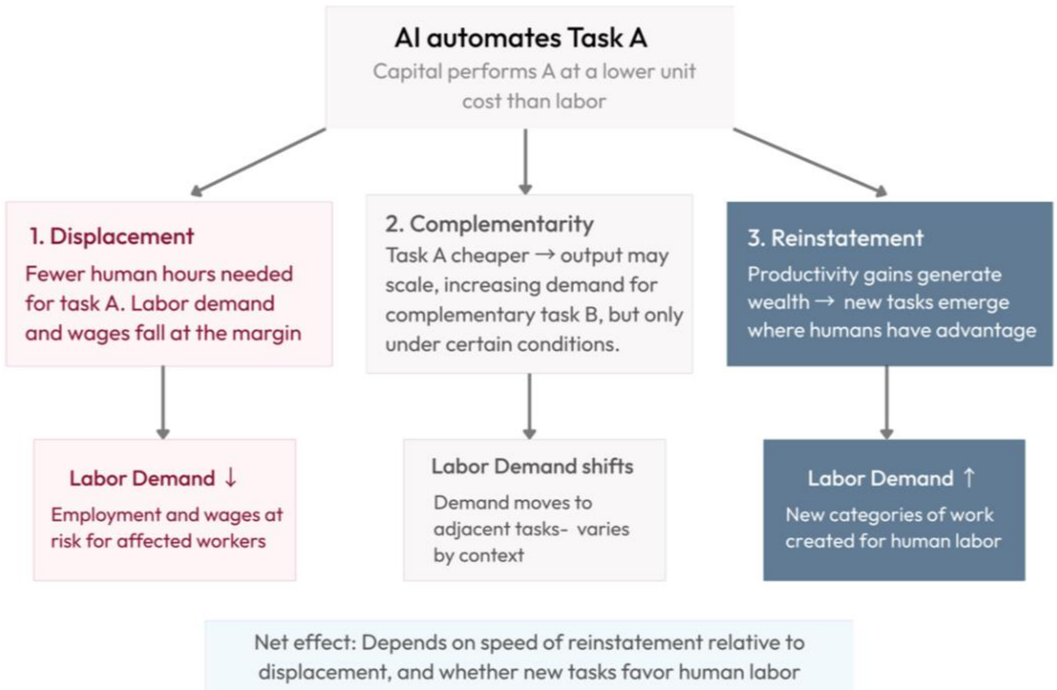

Moreover, the distributional consequences of automation also depend on the nature of the tasks being replaced. Automating routine, codifiable tasks tends to raise the skill floor and can increase productivity in an occupation or industry. Automating expert-level tasks tends to produce the opposite effect: rather than raising the skill floor, it removes the intermediate rungs through which workers build expertise. An MIT study illustrates this dynamic using the digitization of bookkeeping as an example: spreadsheet software raised the skill floor for accountants, increasing their wages, while automating the expert cognitive tasks of inventory clerks made the work more accessible but pushed wages down.[32]

Implications

The central practical implication of the task-level lens is that the effects of AI are already visible inside jobs long before they register as job gains or losses in aggregate employment data. The financial analyst who no longer builds models from scratch but now interprets and pressure-tests AI-generated outputs is performing a different job under the same title. This lag between task transformation and occupational reclassification means that traditional labor market statistics may systematically understate the pace of change.

For workers, the practical implication is that the relevant questions are not "which occupation is safe?" but "which tasks within my occupation are being automated, which are being elevated, and what new skills does the augmented version of my role require?"

[32] David Autor & Neil Thompson, Beyond Job Displacement: How AI Could Reshape the Value of Human Expertise, MIT (2025).

5. Preliminary Labor Market Signals Indicate AI's Pressure on Entry-Level Roles

While the aggregate labor market picture remains healthy, that picture obscures a more complicated story one layer down. Entry-level knowledge work—the traditional on-ramp to white-collar careers—has been contracting in recent years and the pattern has been consistent enough across data sources to warrant serious attention. While overall employment continues to grow, entry level occupations have remained generally stagnant since late 2022, signaling a divergence in the professional lifecycle.[33]

The direction of the effect has not been in dispute. New graduates have been entering a job market that has been largely frozen. Revelio and Bloomberg found a 35% decline in U.S. entry-level job postings since January 2023.[34] Gusto found a 16% year-over-year decline in new grad hiring for 2025, a trend that has been present since 2022.[35] The Burning Glass Institute's analysis of millions of job postings found that the share of roles requiring three years of experience or less dropped sharply between 2018 and 2024 in fields most exposed to AI, while total postings in those fields stayed flat or grew and senior hiring remained stable.[36] Recent employer survey data suggests the downward trend in entry-level hiring may be temporary, but the long-term implications will not be known for some time.[37]

What remains contested is direct attribution. Studies utilizing macro-level survey data, such as the Current Population Survey (CPS), consistently find small or null effects of AI and point to the post-pandemic correction, the Fed rate cycle, and broader economic uncertainty as more plausible primary drivers. One such study observed that, on average, the most AI-exposed occupations showed no substantial difference in employment or earnings growth compared to the least exposed.[38] The Yale Budget Lab identified only a faint divergence in the job mix between recent and established graduates, characterizing the shift as a symptom of a cooling labor market or a post-pandemic correction rather than technological displacement. This

---

[33] The Labor Market for Recent College Graduates: Unemployment, Federal Reserve Bank of New York (last accessed 2026).
[34] Lisa K. Simon, Is AI Responsible for the Rise in Entry-Level Unemployment, Revelio (2025).
[35] Tom Bowen, Graduating into a Slowdown: Class of 2025 Meets a Frozen Job Market, Gusto Insights (2025).
[36] Gad Levanon, et al., No Country for Young Grads: The Structural Forces That are Reshaping Entry-Level Employment, The Burning Glass Institute (2025).
[37] Job Outlook 2026: Spring Update, National Association of Colleges + Employers (2026); The 2025 ZipRecruiter Annual Employer Survey, ZipRecruiter (2025).
[38] Bharat Chandar, Tracking Employment Changes in AI-Exposed Jobs, Stanford Digital Economy Laboratory (2025).

perspective is reinforced by the fact that some effects on entry-level hiring predate the November 2022 public release of ChatGPT.[39]

Studies using more granular data reach different conclusions. A Stanford study using ADP payroll data (2025) identified a 16% relative decline in employment from late 2022 to mid-2025 for workers aged 22–25 in highly exposed occupations, even after controlling for firm-level shocks.[40] Meanwhile, employment for workers aged 35–49 in the same exposed roles grew by over 8% in the same period.[41] The gap is not explained by interest-rate sensitivity, remote-work effects, or technology-sector overhangs—it persists across industries and occupation types. Revelio Labs similarly found that a 10-percentage point increase in AI exposure is associated with an 11% drop in entry-level demand but a 7% rise in non-entry-level demand within the same occupations.[42] While these findings are compelling, the extent of AI's role in the disruption is disputed. The preliminary and proxy data are insufficient to determine a direct causal link and future studies are needed to combine granular payroll data with confirmed firm-level AI adoption data.

Why Entry Level Workers First:

Ultimately, the decline in entry level hiring over the last several years may stem from a combination of AI-driven and cyclical or structural effects that have an outsized impact on early-career positions in the following ways:

AI task substitution. Generative AI disproportionately substitutes for the codifiable tasks that can be articulated as rules and procedures, which make up the bulk of certain entry-level professionals' workload. In contrast, senior workers with a greater pool of expertise and tacit knowledge appear meaningfully less exposed to AI automation. Critically, effects are concentrated in occupations where AI deployment is primarily automative rather than augmentative—where the technology substitutes for human labor rather than complementing it.[43]

Post-pandemic lean staffing. Many organizations discovered during the pandemic that they could sustain productivity with significantly smaller teams. This shift has led to a decoupling of

---

[39] Martha Gimbel, et al. Evaluating the Impact of AI on the Labor Market: Current State of Affairs, The Budget Lab at Yale (2025).
[40] Erik Brynjolfsson, et al., Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence, Stanford Digital Economy Lab (2025).
[41] Erik Brynjolfsson, et al., Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence, Stanford Digital Economy Lab (2025).
[42] Lisa K. Simon, Is AI Responsible for the Rise in Entry-Level Unemployment, Revelio (2025).
[43] Erik Brynjolfsson, et al., Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence, Stanford Digital Economy Lab (2025).

economic growth from white-collar hiring, where sectors experience expansion without a corresponding increase in entry-level headcount.[44]

The lowest friction adjustment margin. Displacing an incumbent worker with AI involves navigating significant institutional frictions, while cutting entry-level hiring serves as the lowest-friction margin of adjustment.[45] This asymmetry is compounded by firms' limited incentives to invest in training less-experienced workers at scale. More broadly, human capital theory suggests firms may underinvest in training workers who are more likely to leave, as they cannot fully capture the returns on that investment.[46]

Macroeconomic and cyclical headwinds. The Federal Reserve's aggressive rate cycle beginning in March 2022—the most rapid monetary tightening in four decades—hit knowledge-work sectors particularly hard and coincided almost exactly with the onset of entry-level hiring declines.[47] Post-pandemic over-hiring in tech, followed by correction, created an additional headwind concentrated in the same occupations most exposed to AI, making attribution especially difficult.

These forces may be mutually reinforcing: a technology that makes entry-level tasks easier to automate, arriving precisely as firms are under pressure to justify each headcount decision, produces a sharper and more durable contraction than any single cause would generate alone. This effect is amplified by broader economic uncertainty, as firms delay entry-level hiring while reassessing how AI may reshape workflows, further suppressing demand at the margin where adjustment is easiest.[48] The full productivity gains and opportunities for these technologies have yet to materialize, and the composition and volume of labor demand may shift as the adjustment becomes clearer.

The Longer-Term Stakes

Regardless of the degree of attribution to AI, if this trend continues it will have far-reaching implications for the future of work and social mobility. The Dallas Fed estimates that even if all employment losses for young, AI-exposed workers were fully counted, they would raise

---

[44] Gad Levanon, et al., No Country for Young Grads: The Structural Forces That are Reshaping Entry-Level Employment, The Burning Glass Institute (2025).
[45] Erik Brynjolfsson, et al., Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence, Stanford Digital Economy Lab (2025).
[46] Briefing to the Task Force, (March 27, 2026); Gary Becker, Human Capital: A Theoretical and Empirical Analysis, University of Chicago Press (1964).
[47] Jenna Ross, Comparing the Speed of U.S. Interest Rate Hikes (1988–2022), World Economic Forum (2022); Board of Governors of the Federal Reserve System, Federal Funds Effective Rate (FEDFUNDS), FRED, Federal Reserve Bank of St. Louis (2026).
[48] Gad Levanon, et al., No Country for Young Grads: The Structural Forces That are Reshaping Entry-Level Employment. The Burning Glass Institute (2025); Tom Bowen, Graduating into a Slowdown: Class of 2025 Meets a Frozen Job Market, Gusto Insights (2025).

aggregate unemployment by only 0.1 percentage points.[49] While these effects are nearly invisible in national statistics, they are simultaneously real for the affected cohort and pose profound long-term structural risks.

By eliminating some of the foundational tasks that traditionally train junior staff, companies may soon find themselves without a viable means to develop the mid-career experts they will inevitably need. This risk is compounded by longstanding underinvestment in workforce training systems, which may be ill-equipped to replace the lost learning-by-doing traditionally embedded in entry-level roles, leaving a permanent gap where mid-career expertise would normally have emerged.[50]

The decline in entry-level hiring in AI-exposed sectors may contribute to a large-scale mismatch, where a surging supply of college graduates projections indicate a surplus of 7 to 11 million more college-educated Americans by 2034 is colliding with a shrinking number of jobs where those credentials were previously the entry ticket.[51] If this trend is left unaddressed, the resulting structural barrier will increasingly affect economic mobility for recent graduates and potentially slow the career progression of certain white-collar professionals.

However, these observations capture early-stage disruptions rather than a fixed outcome and will likely vary by sector. Much of the long-term trajectory will depend on whether efficiency gains lower prices, expand output, or create enough new demand to offset reductions in labor needs. Hiring may strengthen in some industries as certain roles expand and new job categories are created, even as others consolidate.

---

[49] J. Scott Davis, AI Is Simultaneously Aiding and Replacing Workers, Wage Data Suggest, Federal Reserve Bank of Dallas (2026).
[50] Briefing to the Task Force (March 27, 2026).
[51] Matt Sigelman, et al., The Expertise Upheaval, The Burning Glass Institute (2025).

6. The Extent of AI Adoption Will Be One of the Principal Determinants of the Impact on Labor Market Outcomes

AI's near-term economic impact will be determined less by what AI systems can do than by how quickly and broadly firms deploy them in workflows. The labor market responds to actual deployment and utilization patterns, which explains why dramatic capability advances don't immediately translate into employment changes or productivity gains. Today, adoption is growing unevenly across firms, sectors and geographies. Analyzing what drives adoption and the current state of play is necessary to understanding and forecasting AI's impact on the labor market. Identifying the barriers to adoption is also key to ensuring potential gains from AI are achieved.

Challenges of Measuring Adoption

A primary challenge in quantifying the labor market impact of AI is the conflation of technological exposure with actual adoption. Current discourse relies heavily on exposure metrics, which aggregate the theoretical capabilities of AI across specific tasks; however, these figures represent potential utility rather than empirical deployment. To move beyond speculative analysis, research must prioritize usage-based metrics. This distinction is critical: without distinguishing latent capability from realized application, assessments of AI's impact risk misrepresenting the immediacy of labor market outcomes.

The transition toward usage-based analysis is plagued by methodological heterogeneity and significant data fragmentation. Existing frameworks—primarily divided between self-reported surveys and model telemetry—provide only a localized view of the phenomenon. Survey instruments such as the U.S. Census's Business Trends and Outlook Survey are frequently compromised by variation in who is asked (e.g., executives and managers vs. employees), how the question is asked (e.g., implemented into workflows vs. used at all), and the time horizon under inquiry.[52] This, combined with reporting biases typical for survey-based approaches, gives us only a fuzzy understanding of the current state of adoption. Simultaneously, usage indices derived from specific platforms offer high granularity but do not account for the full taxonomy of occupational tasks defined by O*NET or usage of other LLMs.[53]

[52] Business Trends and Outlook Survey, U.S. Census Bureau (2026); Andy Kemp, Frequent Use of AI in the Workplace Continued to Rise in Q4, Gallup (2026); Alexander Bick, et al., The State of Generative AI Adoption in 2025, Federal Reserve Bank of St. Louis (2025).

[53] Maxim Massenkoff et al., Anthropic Economic Index: Understanding AI's effects on the economy, Anthropic (2026).

Adoption Mechanisms

Broadly, adoption in firms occurs via two mechanisms: bottom-up implementation driven by worker experimentation and natural workflow iteration, and top-down implementation directed and enforced by firm leadership. Both mechanisms produce costs and benefits.

Worker-driven adoption allows effective use cases to surface naturally through trial and error, leveraging employees' on-the-ground knowledge of their own workflows and processes. However, without institutional support, this experimentation may not translate into meaningful workflow reorganization.[54] Several factors can stall bottom-up adoption: coordination and congestion problems in the short run, compliance risks, and a late-mover advantage dynamic in which first movers absorb the costs of testing use cases while late movers observe and implement only what works—without bearing the cost of discovery.[55]

Firm-directed adoption is top-down and slower. Firm executives make decisions about where and how to incorporate AI into processes and workflows, and direct workers to adopt AI accordingly. This approach sacrifices the organic discovery benefits of worker-driven experimentation, but gains something equally important: a governing structure.[56] Top-down adoption can resolve coordination problems, incentivize early and consistent use, manage compliance risks, and establish institutional standards that make adoption durable rather than ad hoc.[57] However, mandating AI use alone is insufficient—directives must be accompanied by clear guidelines, training, and defined use cases. Without this scaffolding, forcing adoption is unlikely to produce meaningful changes in how work is actually done.

Recent experimental evidence suggests that the binding constraint on AI adoption is not access to the technology, but the managerial challenge of discovering where it creates value—what researchers call the 'mapping problem'. A 2026 INSEAD field experiment across 515 high-growth startups found that firms given structured information about how other companies had reorganized their workflows around AI discovered 44% more AI use cases and completed 12% more tasks than those without such guidance.[58] Crucially, both groups had identical access to AI tools and technical training—the difference was knowing where and how to deploy them. This suggests that the most effective adoption approach is neither purely top-down nor purely bottom-up, but a deliberate combination: firm leadership providing the structured frameworks and guidelines that help workers search more broadly across their workflows for where AI can create real value.

---

[54] Briefing to the Task Force, (March 27, 2026).
[55] Task Force member at Plenary 1, (March 11, 2026).
[56] Task Force member at Plenary 1, (March 11, 2026).
[57] Task Force staff at Plenary 1, (March 11, 2026).
[58] Kim, Hyunjin, et al., [Mapping AI into Production: A Field Experiment on Firm Performance.] INSEAD Working Paper (2026).

Barriers to Adoption

Barriers to adoption are both structural and cultural. Beyond the barriers that also constrain AI development—such as demand for compute and infrastructure bottlenecks—there are structural barriers unique to adoption such as a workforce that has largely not been trained in AI technology, a lack of real-time data on skills demand that hampers efficient sorting of workers who are most likely to drive implementation of AI, and regulatory issues. However, in the United States, cultural reticence may be as significant a bottleneck as structural barriers—and could be partially responsible for those structural barriers.[59] Whether or not cultural hesitance is justified, it slows down AI adoption.

AI adoption is imperative if we are to see any economic gains from AI. To quote Stanford economist Dr. Erik Brynjolfsson, “Electricity didn’t change the labor market. What people did with electricity changed the labor market.” We can develop the most capable models, yet their true impact remains unrealized until they are thoughtfully deployed and meaningfully utilized across industries. Ultimately, it is human decisions and actions that shape technology’s impacts and entail intended and unintended consequences.[60]

---

[59] Brian Kennedy, et al., How Americans View AI and Its Impact on Human Abilities, Society, Pew Research Center (2025).

[60] Wanda J. Orlikowski, Using Technology and Constituting Structures, Organization Science Vol. 11, No. at 404-428 (2000).

7. Shifting Skill Demand in the AI-Driven Economy

Technological change has long reshaped the relationship between workers' skills and labor market demand, but the pace and scope of AI-driven change may accelerate this dynamic. As tasks are automated, augmented, or newly created, the skills required for many jobs could shift faster than workers—and the institutions that train them—are able to adapt. This raises the possibility of a widening "skills gap": a mismatch between the capabilities workers currently possess and those increasingly demanded by employers, or the exacerbation of existing imbalances in the labor market.

Across the literature, the skills gap is best understood as a set of overlapping mismatches: between existing skills and new task demands, between human and technical capabilities, and between education systems and labor market needs.[61] Measuring this gap remains difficult, without a robust data source to capture the full extent of AI's impact. Instead, researchers rely on proxies such as job postings, employer surveys, and private datasets, though these remain largely inaccessible, making it challenging to fully assess the scale and nature of disruption. As a result, no single headline figure is sufficient to characterize the scale or nature of the skills gap.

Where AI Is Driving Emerging Skills Gaps

Near-term gaps are most likely where AI capabilities are advancing faster than workers and institutions can adapt, particularly in roles defined by routine or data-rich tasks. Only 7% of workers who used AI reported that AI had not changed their skill requirements.[62] This includes not only clerical and administrative work but also a growing set of cognitive tasks in white-collar roles—such as coding, drafting, and summarization—where performance can be standardized.[63] As a result, disruption is extending into higher-skill occupations with significant exposure to AI-enabled task substitution.[64] Though exposure may be higher for more-educated workers, their ability to upskill and adapt might be stronger than among less-educated workers.[65]

At the same time, gaps are emerging in the areas where AI is expanding job demand more quickly than skilled labor supply is being created. New job categories—including AI engineers, AI strategists, industry-specific workflow integrators, data annotators, and AI system

[61] Working to Win: Rebuilding America's Workforce for an Age of Geopolitical Competition, JPMorganChase PolicyCenter & Center for Geopolitics (2025).
[62] Briefing to Task Force, (March 27, 2026).
[63] David Autor, Why Are There Still So Many Jobs? The History and Future of Workplace Automation, Journal of Economic Perspectives at 3–30 (2015).
[64] Rakesh Kochhar, Which U.S. Workers Are More Exposed to AI on Their Jobs?, Pew Research Center (2023).
[65] Sam Manning & Tomás Aguirre, How Adaptable Are American Workers to AI-Induced Job Displacement, National Bureau of Economic Research (2026).

technicians—are contributing to the estimated 640,000 new jobs added.[66] Additionally, physical skills are increasingly needed to match demand for AI-related infrastructure, with an estimated 2.1 million positions in skilled labor trades projected to go unfilled by 2030.[67]

Existing Skills Gaps in the AI-Era Workforce

These emerging pressures build on existing weaknesses in workforce preparation, including gaps in digital literacy, shortages of advanced technical talent, and structural limits in education systems. Digital skills are now required across most jobs, yet access to computer science and AI education remains uneven, contributing to disparities before workers enter the labor market. This uneven access is underpinned by a profound digital literacy floor that remains unaddressed: currently, one-third of the workforce—including a quarter of all managers—lacks the foundational digital skills necessary to begin building specific AI literacy.[68]

Employers already report significant reskilling needs, indicating that many workers are underprepared for AI-driven changes.[69] This is why some countries, notably India[70] and most recently in the United States,[71] have launched introductory AI literacy courses to level the playing field. However, despite the high demand for these resources, there is a heavy foundational concentration in current learning patterns. According to data from Coursera, half of all AI-related enrollments remain strictly limited to introductory courses, suggesting the workforce has yet to move toward the intermediate or advanced implementation skills required for strategic sectors.[72]

Shortages are particularly acute in strategic sectors, including AI, cybersecurity, semiconductors, energy, and advanced manufacturing, where demand for skilled labor continues to outpace supply. These gaps reflect both limited training pipelines and a broader shift toward more technically intensive roles.[73]

At the same time, the pace of change is shortening the “half-life” of skills, making existing knowledge obsolete more quickly while education systems struggle to keep pace. Despite expectations that AI will augment workers and generate new roles, it remains uncertain whether these opportunities will emerge quickly enough to offset displacement, particularly where

---

[66] Te-Ping Chen, The New Jobs Being Created by AI - WSJ, Wall Street Journal (2026).
[67] Wei Xie & Steven Lewis, Building tomorrow’s workforce today, JLL (2026).
[68] Briefing to the Task Force, (March 27, 2026).
[69] The Future of Jobs Report 2025, World Economic Forum (2025).
[70] Yuva AI for All Course, IndiaAI (2025).
[71] US Department of Labor Launches ‘Make America AI-Ready’ Initiative, U.S. Department of Labor (2026)
[72] Briefing to the Task Force, (March 27, 2026).
[73] Tania Babina & Anastassia Fedyk, The Effects of AI on Firms and Workers, Brookings Institution (2025).

innovation is oriented toward automation.[74] As a result, skills gaps reflect not only workforce readiness but also how new technologies are deployed.[75]

Mismatches in Hiring and Labor Market Signals

While not necessarily attributable to AI, these gaps are increasingly visible in hiring patterns in certain sectors. At the entry level, certain firms are reducing junior hiring while increasing demand for experienced workers who can provide oversight and judgment, narrowing pathways into professional roles.[76] This trend persists despite a widespread belief among leadership that AI could effectively bridge the gap for workers who lack years of experience; while 71% of executives acknowledge this potential for technology to compensate for a lack of tenure, only 25% of firms currently plan to provide the specific training necessary to realize that benefit.[77]

A second mismatch lies between credentials and employer needs. As hiring shifts toward demonstrated skills, particularly AI fluency and applied capabilities, traditional degrees may be becoming less reliable signals of job readiness, a dynamic explored further in the following section.[78] Within firms, hiring is also shifting toward technically skilled individual contributors and away from middle management, flattening organizational structures and raising expectations within roles without corresponding changes in training pathways.[79]

These dynamics are compounded by limited visibility into firm-level AI adoption, making it difficult for workers and training providers to anticipate demand. Furthermore, current weaknesses in the job market for highly exposed roles may not be a simple case of technology replacing humans; instead, part of the slowdown appears to be driven by institutional hesitation, as a combination of uncertainties about the long-term effects of AI, economic policy, and federal funding are likely driving many employers to pause hiring in the near-term.[80] As a result, underemployment and persistent talent shortages can coexist across sectors, particularly in areas such as advanced manufacturing, cybersecurity, and AI.[81]

---

[74] Ayush Chopra, et al., The Iceberg Index: Measuring Skills-centered Exposure in the AI Economy (Version 2), arXiv (2025).
[75] Cisco AI Workforce Consortium, AI Workforce Consortium Full Report, Cisco Systems (2025).
[76] Cisco AI Workforce Consortium, AI Workforce Consortium Full Report, Cisco Systems (2025).
[77] Briefing to the Task Force, (March 27, 2026).
[78] Michael Feroli, et al., AI's Impact on Job Growth, J.P. Morgan Global Research (2025).
[79] Matt Sigleman, et al., Beyond the Binary: How Automation and Augmentation Are Combining to Reshape Work, The Burning Glass Institute (2026).
[80] Briefing to the Task Force, (March 27, 2026).
[81] Briefing to the Task Force, (March 27, 2026).

Skills That Will Remain in Demand in an AI-Transformed Labor Market

Taken together, these shifts point to a reorientation of skill demand rather than a simple reduction in work. As routine tasks are automated, demand is moving toward a combination of durable human skills and higher-level technical capabilities. Interpersonal skills—such as communication, collaboration, empathy, and leadership—remain critical, as they underpin many roles and are difficult to replicate with AI. Similarly, skills rooted in expert judgment—such as contextual reasoning, ethical decision-making, and complex problem-solving—continue to require a human role, particularly in settings requiring accountability and ambiguity management.[82]

However, this reliance on judgment is increasingly dependent on a worker's ability to effectively pilot these new tools; the most significant productivity gains and market value will likely accrue to those who can marry deep, decades-long domain expertise with the specific ability to orchestrate and validate AI outputs. Without this ability to collaborate with the technology, even certain forms of high-level human expertise may face devaluation as systems become more capable.[83]

AI literacy is also becoming a baseline requirement, with growing demand for workers who can not only use AI tools but also adapt and oversee them. More broadly, knowledge work is seemingly beginning to shift from execution toward supervision, validation, and orchestration of AI systems, but it remains challenging to identify which skills will become most valuable.

At the same time, many forms of embodied labor remain relatively resilient. Occupations requiring physical dexterity, in-person service, or high-stakes human interaction, including skilled trades, healthcare, and hospitality, face lower substitution risk. In fact, the physical requirements of the AI economy are generating a massive secondary wave of demand for traditional skilled labor. Because the transition requires a vast expansion of physical infrastructure, such as data centers and hardware manufacturing, the need for specialized trades is surging; for instance, a single major technology firm like Google estimated a requirement for 500,000 electricians to meet its infrastructure goals.[84] This reinforces a divide between highly technical, AI-augmented roles and work grounded in physical and social presence.[85]

---

[82] Harry J. Holzer, The AI-Jobs Paradox, Democracy Journal (February 10, 2026).
[83] Briefing to the Task Force, (March 27, 2026).
[84] Briefing to the Task Force, (March 27, 2026).
[85] Konrad Kording & Ioana Marinescu, Artificial) Intelligence Saturation and the Future of Work, Brookings Center on Regulation and Markets Working Paper (2025).

8. Traditional Educational Pathways and Trainings May Fall Short in Meeting the AI Moment

The modern education and workforce development system is structured around several core pathways: K–12 education, career and technical education, higher education, workforce training programs, and work-based learning.[86] Together, these systems are designed to prepare individuals for long careers across a wide range of occupations, from corporate roles to skilled trades. Each pathway plays a distinct role, from building foundational skills to supporting specialization and reskilling. Work-based learning, in particular, has historically enabled workers to adapt over time by accumulating experience and transitioning into adjacent or more advanced roles.

However, as the structure of work shifts in response to artificial intelligence, the effectiveness of these pathways is increasingly in question. As discussed in Section 4, certain early-career roles that historically functioned as training grounds have become less available, weakening traditional transitions from education into employment. This section focuses on a related issue: whether existing educational pathways are sufficient to prepare workers for an AI-driven labor market, how they may need to evolve, and whether new pathways are required.

At a fundamental level, this challenge reflects a growing mismatch between how education systems are structured and how skills are now developed, deployed, and updated in the labor market. Educational systems are designed to reduce skills gaps by equipping workers with competencies aligned to employer demand, but this function depends on relatively stable skill requirements. Technological change has always reshaped skill demand, but AI may accelerate both the pace and scope of this process, requiring more continuous adaptation.[87]

The Limits of Existing Pathways

Current education systems assume that early-life learning provides stable, long-lasting skills for a career before entering the workforce, and rely on those same competencies over time. This model is breaking down as skills evolve at rapid speed.[88] Underlying skills requirements shift even within some of the most stable occupations, making continuous learning a more critical necessity.[89]

---

[86] Emily Musil, et al., The Computing Imperative: Building America's Talent Engine in the Age of AI, Milken Institute (2025).
[87] David Autor, Why Are There Still So Many Jobs? The History and Future of Workplace Automation, Journal of Economic Perspectives at 3–30 (2015).
[88] The Future of Jobs Report 2025, World Economic Forum (2025).
[89] Daron Acemoglu & Pascual Restrepo, Automation and New Tasks: How Technology Displaces and Reinstates Labor, Journal of Economic Perspectives at 3–30 (2019).

At the same time, traditional credentials are becoming less reliable signals of capability in certain areas. Their value depends on the alignment with labor market demands, but fragmented and outdated data systems make it difficult to track outcomes and the infrastructure to verify this alignment cannot capture real-time transitions across occupations. This has contributed to the creation of the “invisible graduate flow,” in which it is difficult to determine how individuals, particularly those from generalist backgrounds, transition into emerging roles.[90]

These limitations are compounded by broader structural challenges. In certain areas, educational institutions, workforce programs, and employers may work separately from each other, resulting in a disconnection between curriculum design and evolving job requirements. The inability to keep pace with the speed of AI development and adoption contribute to the persistent gaps between training and labor market demand.[91] Finally, the growing disparity between the number of graduates and available high-skill roles contributes to underemployment and credential inflation, further weakening degrees as signals of job readiness while overlooking workers with relevant skills but nontraditional backgrounds.

Expanding and Adapting Existing Pathways

One response is to modify existing pathways to better align with evolving skill demands. This includes further embedding AI literacy across education systems, strengthening connections between education providers and employers, and integrating more applied, work-based learning into curricula.

There is growing consensus that AI literacy is becoming a baseline requirement across occupations, analogous to digital literacy in earlier technological transitions.[92] Efforts to integrate AI concepts into K–12 education and career and technical programs reflect this shift, with the goal of ensuring that workers enter the labor market with a foundational ability to interact with AI systems.[93]

At the same time, there is increasing emphasis on aligning curricula with labor market needs through closer employer engagement. Sector-based training programs, apprenticeships, and community college partnerships are frequently cited as mechanisms for improving this

[90] Briefing to the Task Force (March 27, 2026).
[91] Gad Levanon, et. al., No Country for Young Grads: The Structural Forces That are Reshaping Entry-Level Employment, The Burning Glass Institute (2025).
[92] Emily Musil, et al., The Computing Imperative: Building America’s Talent Engine in the Age of AI, Milken Institute (2025).
[93] Briefing to the Task Force. (March 27, 2026).

alignment.[94] However, scaling these approaches remains challenging, particularly given institutional constraints and uneven incentives across stakeholders.

International models provide useful comparisons. The German vocational system integrates education and work-based training earlier in the lifecycle, enabling students to develop occupation-specific skills before entering the labor market.[95] While effective in aligning training with employer demand, it has been criticized for requiring early specialization, potentially limiting flexibility. This highlights a broader tension between early alignment and long-term adaptability.

The Emergence of Alternative Pathways

Beyond modifying existing systems, new pathways are emerging that operate alongside or outside traditional education structures. We highlight four approaches that reflect different ways of acquiring and signaling skills in a more dynamic labor market.

One category includes modular and stackable credentials, such as professional certifications and micro-credentials, that allow workers to acquire targeted skills over time. These credentials can serve as alternative signals of competency, particularly for mid-career reskilling; however, their value depends on consistent employer recognition.[96]

A second category includes employer-led training and work-based learning. Firms are increasingly investing in internal programs that embed learning within workflows, blurring the boundaries between education and work.[97] However, these programs can be resource-intensive, limiting employer uptake, and may produce skills that are not easily transferable across employers.

A third category includes online platform-based learning ecosystems, which provide scalable access to rapidly evolving technical skills. Enrollment data suggest strong demand for foundational AI skills: Coursera alone has recorded over 16.8 million global enrollments in generative AI courses, with rapid growth in recent years.[98] At the same time, learners are prioritizing short, applied courses that can be integrated into existing work schedules, reflecting

---

[94] Harry J. Holzer, Understanding the Impact of Automation on Workers, Jobs, and Wages, Brookings Institution (2022).
[95] A Guide to the German School & Education System, CBS University of Applied Sciences (2026).
[96] Alex Swartsel, et al., The AI-Ready Workforce: How Leaders and Workers Can Prepare for a Reshaped Future of Work, Jobs for the Future (JFF) and Intel (2023).
[97] Cisco AI Workforce Consortium, AI Workforce Consortium Full Report, Cisco Systems (2025).
[98] Briefing to the Task Force, (March 27, 2026).

a shift toward continuous, career-driven learning. Still, cost barriers and uneven course quality and returns for credentials remain great challenges.[99]

Finally, AI itself is emerging as a lower-cost training mechanism, providing a new style of more personalized and conversational instruction, offering real-time feedback and accelerating skill development particularly for less experienced workers.[100]

Coordination, Credibility, and System-Level Implications

The expansion of alternative pathways has increased the number of ways workers can acquire skills, but not the system's ability to translate those skills into clear labor market signals. The core challenge is not a lack of pathways, but the absence of a coherent system that connects them.

Disconnection between credentials, training programs, and learning platforms result in inconsistent skill interpretation, weakening the value of credentials and perpetuating reliance on traditional degrees, particularly among smaller firms. Simultaneously, limited outcome data hinders the assessment of effective pathways, leaving workers and institutions with incomplete information. These fragmentations come to impede participation due to costly, time-intensive, and difficult-to-evaluate navigation. This suggests the future of education pathways will require integrating multiple training programs into a continuous, nonlinear learning system embedded within the labor market itself.

In practice, this shift implies several areas of focus. First, improving the applicability of credentials will be critical to restoring their ability to signal job-readiness. Second, establishing responsive data infrastructure will allow for tracking skill acquisition and outcomes, enabling better informed decision-making by institutions. Third, stronger coordination throughout the workforce pipeline is necessary to align training with evolving demands in real time.

At the same time, expanding access to these pathways will require reducing barriers related to cost, time, and information. This includes not only increasing the availability of training, but also ensuring that workers can engage with it without significant disruption to employment or income.

Ultimately, the challenge is not simply to expand education pathways, but to make them function as a system—one that is legible to employers, accessible to workers, and capable of adapting alongside technological change.

---

[99] Matt Sigelman, et al., Holding New Credentials Accountable for Outcomes: We Need Evidence- Based Funding Models, Burning Glass Institute

[100] Erik Brynjolfsson, et al., Generative AI at Work, The Quarterly Journal of Economics at 1–48 (2025).

9. The Distribution of AI's Gains Need Not Be Concentrated, and Can Be Shaped by Implementation Choices

The relationship between AI and economic inequality is not predetermined. AI could reduce wage inequality by extending high-productivity tools to those who previously lacked access, or it could accelerate the concentration of gains among those already well-positioned. Inequality in the modern world largely results from the uneven dissemination and adoption of technologies, a process that is dictated by whether a society's institutions are inclusive or extractive.[101] Democratizing access to AI will break down barriers to help ensure these tools are accessible to all rather than limited to a few. Ultimately, which outcome prevails will depend less on the technology itself than on the choices made by firms, policymakers, and workers in the years immediately ahead.

The optimistic case is grounded in real evidence. Research on earlier information technologies suggests that tools that augment cognitive tasks can raise the productivity and, over time, the wages, of workers who might otherwise be left behind by automation.[102] When AI assists a mid-career worker in drafting documents, analyzing data, or navigating complex regulatory environments, it can reduce the skill premium that previously made those tasks the exclusive domain of highly credentialed professionals. Several recent studies examining the deployment of AI-assisted tools in customer service, legal research, and clinical settings have found that less-experienced workers often benefit disproportionately from AI assistance, narrowing within-firm performance gaps rather than widening them.[103] If these patterns hold at scale, AI could function as a genuine equalizer; not by eliminating skill distinctions, but by lowering the cost of developing competence.

The mechanism described above cuts both ways—AI-driven shifts in skill demand also could lead to reduced demand for labor in certain occupations, displacing workers in the process. Reabsorption into the workforce is not instantaneous: even where AI lowers barriers to entry in new or adjacent occupations, workers require time to retrain, relocate, and signal new competencies to employers. The empirical record on prior transitions—most rigorously documented in studies of trade-induced displacement—suggests this adjustment period spans years, and that the burden falls unevenly, concentrated among workers with less geographic

---

[101] Daron Acemoglu & James A. Robinson, Why Nations Fail: The Origins of Power, Prosperity, and Poverty, Crown Currency at 53 (2012).

[102] David H. Autor, Why Are There Still So Many Jobs? The History and Future of Workplace Automation, Journal of Economic Perspectives (2015).

[103] Erik Brynjolfsson, et al., Generative AI at Work, The Quarterly Journal of Economics at 1-48 (2025).

mobility, fewer financial resources to weather unemployment, and less access to retraining infrastructure. [104]

Further, AI could more fundamentally reshape the value of labor in general. Since the Industrial Revolution, labor has been the main bottleneck in production. Its relative scarcity is what made it valuable—capital is reproducible, labor is not, and that asymmetry underpinned a century of broadly rising wages. AI disrupts this asymmetry by enabling capital to substitute for the cognitive labor input itself.[105] If AI absorbs more work than it generates, the result is an economy that grows without broadly sharing that growth—then labor's share of economic output contracts. The economy could grow significantly while the gains flow predominantly to those who own the capital doing the work. That outcome is not inevitable. The outcome in which AI raises living standards broadly while compressing rather than widening inequality is achievable, but it is not necessarily the default, and it may not arrive on its own.

Ensuring this outcome will require getting a number of things right—deliberately and soon—on access to AI tools, on how firms structure and allocate human-AI workflows, on investment in education and training, and on the institutional frameworks governing how productivity gains are distributed.

Even at the firm level, adoption creates winners and losers. Firms that adopt AI fastest tend to be larger and better-capitalized. Workers in those firms are more likely to receive training, to have job functions that lend themselves to human-AI collaboration, and to capture a share of the productivity gains through wages or advancement. Workers in smaller firms or in occupations with thinner margins and less institutional capacity for technology adoption may find themselves on the wrong side of an AI productivity gap, not because their work is being automated away, but because they are not gaining access to the augmentation benefits their peers elsewhere enjoy.[106]

This advantage gap also has a significant geographic component—firms adopting AI tend to be concentrated in metropolitan areas with existing talent advantages; similarly, AI development occurs in a small number of technology hubs. Workers in rural or economically distressed areas are less likely to receive the skills training firms adopting AI provide. Advantages accruing to dominant regions are therefore self-reinforcing. Concentration draws in further business investment, talent, and infrastructure, creating a gap that is increasingly difficult to close.[107] The communities most exposed to displacement risk from AI, by contrast, tend to be those with

---

[104] David H. Autor, et al., The China Shock: Learning from Labor Market Adjustment to Large Changes in Trade, Annual Review of Economics (2016).
[105] Briefing to the Task Force, (March 27, 2026).
[106] Stefan Koopman, et al., The Economic Impact of AI: Four Scenarios, Rabobank (2024).
[107] Sandrine Kergroach & Julien Héritier, Emerging Divides in the Transition to Artificial Intelligence, OECD Publishing (2025).

higher concentrations of routine-task employment, limited access to retraining infrastructure, and fewer institutional anchors capable of managing a significant labor market transition. History offers cautionary precedent here: the gains from prior technological transitions were real, but they were not broadly shared in time or place, and the adjustment costs fell heavily on specific workers, families, regions, and institutions that lacked the resources to adapt.

None of this argues for slowing AI adoption. The productivity gains AI makes possible are among the most promising mechanisms available for raising living standards broadly, and forgoing them would impose its own costs, particularly in the context of national competitiveness. But productivity gains that flow narrowly will not serve the nation's long-term interests, economic or otherwise. A workforce in which a significant fraction of workers feels that technological progress is something that happens *to* them rather than *for* them is not a foundation for durable competitiveness or social cohesion.

The implication is that inequality must be treated as a first-order design problem for AI policy, *not* an afterthought to be addressed once the technology has matured. Decisions made now about access to AI tools in education and training, how firms structure and distribute human-AI workflows, which communities receive investment in digital infrastructure, and how the gains of AI-enabled productivity are shared will shape the distributional trajectory of this technology for a generation. The window for making those decisions well is open, but it is not unlimited.

10. The Need for Coordinated, Whole-of-Nation Action

The evidence assembled in these findings leads to a single, unavoidable conclusion: the United States should not rely on any one sector, institution, or level of government to navigate this transition alone. The scale of AI's impact on work—its speed, its unevenness, and its capacity to either broadly distribute or narrowly concentrate economic gains—demands a response that is equally broad. What is required is a coordinated, whole-of-nation strategy that aligns the interests of workers, employers, educators, and policymakers at each critical inflection point along the path from AI capability to labor market outcome. The decisions made in the next several years will determine whether that strategy exists, and whether it arrives in time.

The effects of AI on the future of work will ultimately reflect the cumulative actions and decisions of individuals and institutions across the country. That is both a warning and an opportunity. Where interests are misaligned, the transition will be painful and uneven. Employers who face no incentive to reskill, workers who lack access to retraining, and educators preparing students for a labor market that no longer exists are not edge cases; they are the default outcome in the absence of coordinated action. Where those interests are deliberately aligned, the outcome can be different. The whole-of-nation approach this Task Force recommends is not an abstract aspiration; it is the practical recognition that mitigating uneven impacts across individuals, firms, industries, and regions requires coordination that no single actor can provide alone.

At the task level, the main dependencies revolve around the interplay between the capabilities and adoption of AI tools. AI capabilities and the infrastructure to support them will be key factors in the scope of automation in work tasks.[108] As these capabilities continue to change, they will also determine which human skills remain complementary to AI-based work.[109] However, the capabilities alone will not influence work tasks without adoption. Worker-driven and employer-driven AI adoption pathways will be key determinants of whether tasks are more augmented or automated, and the timeframe over which workflows substantially shift.[110]

---

[108] Pascual Restrepo, We Won't Be Missed: Work and Growth in the AGI World, National Bureau of Economic Research (2025).
[109] Erik Brynjolfsson & Tom Mitchell, What can machine learning do? Workforce implications, Science (2017).
[110] The AI Ready Workforce, Jobs for the Future (2023).

As task changes expand out to the workforce level, the economics of AI adoption will start to play a larger role. The interaction between tech, labor, employers, and educators will be key to understanding the extent to which hiring and investment decisions change as task-level adoption of AI changes. Worker replacement becomes more likely if AI is cheaper than equivalent human labor, and if that human labor is not reinvested and reskilled into new value-add roles.[111]

Finally, macroeconomic effects will depend on the breadth of changes and how policies adapt. Workforce disruption that is narrow in scope or confined to certain industries would have limited macroeconomic fallout if workers are able to shift to other growing or emerging fields.[112] Responses must harness the upside potential of AI in order to stay globally competitive relative to other countries.[113]

The United States has navigated technological transitions before, and it has risen to meet challenges that once seemed insurmountable. This moment is no different in kind, but it may be greater in scale and speed than anything that has come before. The country enters this transition with something it lacked in prior ones: advance warning and a clearer picture of where the risks are concentrated than has ever been available at this stage of a major technological shift. That is cause for confidence, not complacency. The United States has the tools, the talent, and the institutional capacity to lead this transition.

---

[111] Erik Brynjolfsson & Tom Mitchell, What can machine learning do? Workforce implications, Science (2017).
[112] Daron Acemoglu & Pascual Restrepo, Automation and New Tasks: How Technology Displaces and Reinstates Labor, Journal of Economic Perspectives (2019).
[113] Erik Brynjolfsson, et al., A Research Agenda for the Economics of Transformative AI, National Bureau of Economic Research (2025).

Appendix 1. Task Force Staff and Contributors

SCSP Task Force Staff

Dr. Ryan Carpenter
Veronica Jijon
Dr. Nandita Balakrishnan
Carina Ritcheson
Caroline Armstrong
James Ryseff
Brock Dodds
Sara Huta
Ylber Bajraktari

NVIDIA Task Force Staff

David Shahoulian
Meg King
Faiza Khan
Sarah Weinstein
Sophie Goguichvili

Advisors to the Task Force

Dr. Harry Holzer
Clara Kaluderovic
Alex Kotran
Maria Flynn
Dr. Diana Gelhaus