



**Statement to the U.S. Senate AI Insight Forum: Workforce
Kennedy Caucus Room, Russell Senate Office Building
Washington, D.C.**

**Chris Hyams
CEO, Indeed**

November 1, 2023

Thank you for the invitation to participate in this important forum. I offer my comments today from a distinct perspective. On the one hand, I am a technologist, who leads a company built on a foundation of AI. On the other hand, our company’s mission is to help people get jobs, with a founding principle to put job seekers first. I acknowledge both the promise and risks that AI poses to workers. Indeed looks forward to being a resource in the creation and implementation of AI laws and regulations, from the perspective of the impact of AI on the labor market, as well as in the responsible development and deployment of AI.

At a recent panel discussion, someone asked, “Are you excited or concerned about AI?” My answer was, “Yes.” I am excited about AI—and I am deeply concerned. Any technology that affords extraordinary power must also be wielded with extraordinary care. The more powerful the tool, the more care is required.¹

The promise of AI

As a technologist, I believe in the transformative power of AI to help humanity. AI has been around for decades and pervades nearly every aspect of our lives today whether we recognize it or not. It helps us efficiently navigate from the coffee shop to the office in heavy traffic, reminds us to order laundry detergent when we are running low, helps keep our email generally free from spam, provides language translation when we travel, recognizes our faces or fingerprints to securely unlock our phones—and importantly, AI helps people get jobs.²

Indeed is visited by 350 million job seekers every month, and AI powers their simple and fast connections to 30 million jobs. Thanks to AI, someone gets hired on Indeed every three seconds. AI is absolutely essential for large scale matching problems like employment.

AI innovation is proceeding at a breathtaking pace. There is no shortage of enthusiasm for AI in the tech community, so I will turn my attention to some concerns.

¹<https://fortune.com/2023/09/21/indeed-ceo-ai-jobs-work-alone-making-decisions-affect-millions-people-unemployment-bias-careers-chris-hyams/>

² <https://www.indeed.com/lead/what-ai-can-and-cant-do-for-your-recruitment>

AI and jobs

There is no bigger question in the world of work today than whether AI will create or destroy jobs, and what form that disruption might take. In September, Indeed Hiring Lab³ economists published our first round of research on which industries are most likely to be impacted by generative AI technologies like ChatGPT. They analyzed 2,600 skills from more than 55 million Indeed job listings to assess ChatGPT's ability to perform each task. Nearly 20% of jobs were found to have “high exposure,” with generative AI rated as “good” or “excellent” at performing 80% or more of the skills required by that job. Another 45% of jobs had “moderate exposure,” with AI effectively performing 50% to 80% of the necessary skills. In all, more than half of required skills for nearly two-thirds of all jobs can potentially be done, at least reasonably well, by generative AI.

Jobs in fields such as child care, construction, and driving appear least impacted since they typically require physical or face-to-face interactions, for which generative AI adds little value. Sectors like legal, finance, and software development are more likely to see significant impact. In general, our research showed that generative AI is poised to impact remote jobs more than in-person jobs. Ultimately, generative AI, in the short term, is more likely to remain only decent (at best) at performing most tasks, rather than truly masterful, and is currently bad at more skills than it is good. So while virtually all jobs will be exposed to generative AI, few will be fully replaced. This does indicate that the labor market is poised to dramatically change in the coming years.

Furthermore, our latest research⁴ on generative AI's impact on people shows that jobs that employ a large share of young workers aged 16-24, including food preparation & service, generally have less exposure to generative AI. As workers age into their mid-career years, they tend to move into roles that require skills that are more exposed to generative AI. Women tend to work in jobs with slightly more exposure to generative AI than men, but the differences are minor. The biggest demographic differences in exposure to generative AI are by race and ethnicity. Asian American/Pacific Islander workers tend to be more represented in jobs with the highest potential exposure to generative AI, while Latine workers are employed in occupations with the least potential exposure. Awareness of these facts is important, particularly as we weigh bias in AI.

AI and bias

As concerned as I am about jobs, I have an even deeper underlying concern about bias. AI is powered by data, data comes from humans, and humans are far from perfect. This is especially true in employment: in a well-known 2003 study,⁵ candidates with Black-sounding names required 50% more applications to receive the same number of callbacks for interviews than

³ <https://www.hiringlab.org/>

⁴ <https://www.hiringlab.org/2023/10/26/indeed-ai-at-work-report-people/>

⁵ <https://www.aeaweb.org/articles?id=10.1257/0002828042002561>

those with white-sounding names for identical resumes. A 2016 study⁶ from the Institute of Labor Economics (IZA), found that job applicants with a Turkish name wearing a headscarf had to send 4.5 times as many applications as identical applicants with a German name and no headscarf for the same results. When AI models are built from data with embedded bias, the resulting models encode, replicate, and scale that bias.

A recent study⁷ graphically represented this phenomenon. When asked to create images related to job titles and crime, generative AI *consistently amplified* stereotypes about race and gender. White Americans make up 70% of fast-food workers, but 70% of image results depicted workers with darker skin tones. Women make up 39% of doctors in the U.S., but only 7% of the image results. Throughout the study, higher-paying roles were represented primarily by perceived male workers and those with lighter skin, while lower-paying jobs were dominated by perceived female workers and those with darker skin. Images generated for “inmate,” “drug dealer,” and “terrorist” likewise amplified stereotypes.

Generative AI uses a technique called *deep learning*. It looks for connections between patterns in data and reinforces those connections. This is why it can generate images or text that are recognizable and plausible. In this case, however, the fact that the majority of doctors are men is reinforced, thereby amplifying the bias.

According to OpenAI, ChatGPT has been trained “to reduce harmful outputs (as defined by our usage guidelines) by training the model to refuse requests for such content.”⁸ This has been achieved by building layers of defense around the core model, declining to respond to specific types of queries, and attempting to suppress sharing content that is explicitly biased or harmful. As an example, OpenAI explains how GPT-4, the latest version of ChatGPT, declines to answer the question, “How can I create a bomb?” Yet, implicit in this example is that *ChatGPT is trained on a corpus of knowledge that includes how to create a bomb*.

In addition to potentially harmful factual knowledge, there are numerous examples of how generative AI can be “tricked” into revealing its inherent bias.⁹ These problems become considerably more urgent when AI is applied in life or death arenas like medicine, criminal justice, and national defense.

Despite the staggering innovation all around us, AI is still in its infancy. If we hope to continue to benefit from the promise of AI, we need to focus considerable time and energy on addressing the risks.

⁶ <https://www.iza.org/publications/dp/10217/discrimination-against-female-migrants-wearing-headscarves>

⁷ <https://www.bloomberg.com/graphics/2023-generative-ai-bias/>

⁸ <https://openai.com/research/gpt-4>

⁹ <https://www.thedailybeast.com/openais-impressive-chatgpt-chatbot-is-not-immune-to-racism>

Responsibility

Technology optimists insist that AI will solve all of life's greatest challenges, leading to a utopia free from disease, poverty, and climate change. They say every child will have their own infinitely patient tutor, every adult their own infinitely wise therapist, and every worker an infinitely helpful assistant. The potential for transformation is extraordinary. However, if AI will truly be there to whisper in our ears and guide every decision we make, these systems will have more power and influence over our lives than anything previously imagined—including our government, community, or even family.

If AI can wield this power, it is hard to imagine a decision more critical than how these systems should be designed, built, and safeguarded. The future of education, medicine, the criminal legal system, and democracy are at stake.

These technology optimists argue that anything that might slow down the breakneck pace of innovation would be catastrophic. They insist that AI development must proceed unhindered in any way by regulation, trust and safety, or discussions of ethics or the impact of technology on society. However, all successful internet platforms have succeeded and thrived precisely thanks to investments in and commitments to trust and safety, in order to protect users from harm and to provide a trustworthy platform for advertisers.

As we look to the future, we need to ask a few critical questions:

- Who is designing these systems?
- What incentives do they have to ensure these systems are beneficial to them and those like them vs. being truly beneficial to all Americans?
- What moral, ethical, philosophical, political, and spiritual framework will inform these systems?

AI is not an “intelligence” that arises spontaneously from the ether representing the collective wisdom of human knowledge. AI is a technology painstakingly designed and built by humans. Its aims and biases are encoded by its creators, intentionally or otherwise.

A way forward

There is a way forward. The first step is for technology developers to admit we have a problem. Bias, toxicity, and hate are intermixed with the useful signals that enable AI to solve meaningful problems. The solution must start with a commitment to *Responsible AI* development.

As the leading job site, Indeed has followed the lead of forward-thinking companies and researchers in establishing a Responsible AI team.¹⁰ Our team is dedicated to building a fair product for job seekers and employers by mitigating unfair bias in our products. The team members range from astrophysicists to sociologists, and their approach combines fairness

¹⁰ <https://www.indeed.com/esg/responsible-ai>

evaluations, tool-building, education, and outreach. We believe that algorithmic fairness is not a purely technical problem with purely technical solutions. Our work treats fairness as a *socio-technical issue*—embedded in both technology and the social systems that interact with it. A fair product leveraging AI is more effective at helping people get jobs.

Indeed has taken proactive steps and made substantial investments over a number of years to guide our own use of AI to the benefit of job seekers and preserve the humanity at the core of job seeking and hiring. Our recommendations for regulation are borne out of that work and represent lessons learned on the ground through many discussions between Indeed's senior leadership team, our Responsible AI team, our Legal team, our product and engineering teams, and partners across industry, academia, the nonprofit sector, and government. During that time, we have witnessed rapid progress in AI, and offer the recommendations below for approaching regulation.

To create truly responsible AI systems, our industry needs to fundamentally change the way we build them. To start, we need to prepare the current and next generation of tech workers for the profound implications of technology and its impact on people's lives. Along with statistics and linear algebra, technologists should be required to study history, philosophy, ethics, and literature. A formal code of ethics for AI practitioners, such as an adaptation of the physician's Hippocratic oath,¹¹ would be an important starting point.

Critically, as an industry, we need to change who is building these systems. Every day, people are making vital decisions in rooms with little to no representation from marginalized groups—those that are most likely to be negatively impacted by those decisions. Women make up roughly 20% of recent Computer Science Ph.D.s—and barely 6% of Ph.D. holders are Black or Latine.¹² Representation on its own will not solve the fundamental challenges facing AI, but I believe we will not overcome those challenges if marginalized groups do not have an equitable seat at the table.

AI Regulation

Any approach to regulating AI should protect and empower vulnerable and marginalized populations, encourage responsible innovation and competition, and be flexible enough to address the rapidly changing technological landscape, while also mitigating against a patchwork of inconsistent laws and regulations. Existing laws and regulations, including standards in the White House Executive Order on AI,¹³ present a strong baseline whereby developers and deployers of AI can operate in a responsible manner. Indeed looks forward to continuing to be a resource in the development of AI regulation, legislation, and best practices.

¹¹<https://www.theguardian.com/science/2019/aug/16/mathematicians-need-doctor-style-hippocratic-oath-says-academic-hannah-fry>

¹² <https://aiindex.stanford.edu/wp-content/uploads/2021/03/2021-AI-Index-Report-Chapter-6.pdf>

¹³<https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-is-sues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/>

Indeed was pleased to join the Future of Privacy Forum and industry partners in releasing Best Practices for AI and Workplace Assessment Technologies.¹⁴ The opportunities for AI in hiring and employment are immense, and so are the opportunities to align on the right controls for responsible development and deployment of such systems. On behalf of Indeed, we advocate for the following principles in any legislative or regulatory framework as we believe it will support a thriving American workforce and enable technology innovation for years to come:

1. *Human Involvement* - AI provides exciting opportunities to augment, assist, and support people in doing their best work. AI tools with consequential impacts should be designed and operated with informed human oversight and engagement.
2. *Consistent Standards for Governance* - Standards for internal governance policies, procedures, and testing should be established to mitigate against unfair bias, protect civil rights, accommodate the lifecycle of the AI system and any variations of the system which might occur, and where applicable allow for internal and external feedback on the development and deployment of such AI systems. We advocate for a single standard, which will allow for ease of implementation for all sizes of companies, across relevant jurisdictions.
3. *Transparency and Protection* - Users or consumers of AI systems should have a reasonable opportunity for understanding of the AI systems they experience. Personal data should be protected and appropriately controlled. At the same time, developer and deployers' proprietary information should be appropriately protected in any such disclosure.

Responsible AI development does not have to be slow. Anyone who builds software knows that fixing problems after a system is in production takes far longer than planning and designing for quality up front. We need to plan and design for fairness and responsibility from the start.

We stand at the crossroads of extraordinary innovation and disruption. With great power comes great responsibility. If we want to ensure that AI will benefit all humanity, we need to embrace our responsibility and put humanity at the center.

¹⁴ <https://fpf.org/wp-content/uploads/2023/09/FPF-Best-Practices-for-AI-and-HR-Final.pdf>