

Humans in the Loop: The evolution of work in early experiments with Generative AI

Executive Summary



April 2026

Summary findings from the
MIT Working Group on Generative AI and the Work of the Future



Authors
Ben Armstrong
Julie Shah

Contributors
Kate Kellogg
Sabiyah Ali
Greer Brigham
Carey Goldberg

Shakked Noy
Purna Ravi
Azfar Sulaiman
Leana Tejedor

Felix Wang
Whitney Zhang

HOW CAN AI LEAD TO BETTER WORK? LESSONS FROM EXPERIMENTS AT LEADING COMPANIES.

For three years, a group of researchers at MIT has been studying how companies have deployed generative AI, and how work is changing as a result. Drawing on interviews and workshops with more than fifty companies, they found that companies across various industries – healthcare, insurance, finance, manufacturing, and retail – all seemed to point generative AI applications at three common problems.

They tried to use AI to speed up annoying bottlenecks that prevented workers from focusing on the tasks that they valued most in their jobs. They used AI to synthesize what multiple domain experts might say about a topic so one individual could take on a project that might have previously required a diverse team. And they sought for AI to help novice workers climb the learning curve in a new domain, gaining capabilities quickly that might have otherwise taken years to acquire.

Early evidence from employers' experiments trying to solve these problems with AI suggests six lessons for how to capture the benefits of these technologies for productivity and knowledge creation without reducing the quality of work or the level of skill among workers who come to rely on AI.

Minimize drudgery: Although some workers in [surveys](#) report enjoying their routine tasks, there are many workers who would prefer problem solving and creative tasks to the mundane and routine. Some organizations have ventured into applications of generative AI that aim to augment tasks that workers enjoy – and from which they derive value – setting up resistance from the workers affected. The more straightforward applications of generative AI can focus on the bottlenecks that both workers and employers can agree would be better if they were gone.

Promote learning: There is [early evidence](#) that using generative AI can lead to forms of mental offloading, where a worker performs a task using generative AI, but does not retain the knowledge of the task that they performed. Learning at work can unlock better career opportunities for employees, and deliver higher productivity for employers. There is evidence that generative AI technologies can help employees learn by pinpointing where there are gaps in their knowledge and providing relevant information. However, ensuring that generative AI tools are used for learning and not for offloading will require guardrails that steer workers in the former direction.

Preserve teamwork: Generative AI applications can decrease individuals' reliance on others for expert advice. Even in the cases where generative AI tools can enable an individual to do the work of a team, it may be important to preserve teamwork for other reasons, including mentorship, collective learning, and trust building among people who might rely on one another for other tasks.

Interface design: As employers have grown to understand the capabilities of generative AI tools, they have tended to buy tools from vendors rather than build their own custom applications. Even when organizations use tools built for their industry, they have a choice in how their workers engage with generative AI: they can customize the interfaces that their workers use. Good interface design can ensure that generative AI promotes learning and attention rather than skill atrophy. The principles for good interface design have sought to maximize a user's situational awareness – their ability to understand what is happening and why – as well as to manage their mental workload.

Domain expertise: There is some indication that generative AI tools can help users extend their capabilities. However, this should not be mistaken as a replacement for domain expertise. The users of AI tools must still be able to provide the underlying context, interpret their results, and improve them as they improve their understanding of the background domain. In short, domain expertise can still complement generative AI technologies, but the expertise itself must be sufficient to validate and scrutinize what generative AI tools provide.

Accountability: Generative AI tools can create a moral hazard for workers using them. A worker may have an incentive to use generative AI tools as a shortcut to produce information and perform work that appears valid and impressive, even if that work contains errors or masks a lack of underlying understanding or capability. One way for organizations to address the moral hazard challenge is to establish accountability for workers using AI tools to increase their incentives for learning and raise their costs of making an error. Employers can introduce random errors or tests of their employees to audit performance and reward workers for their vigilance.