

AI TOOLS IN THE WORKPLACE: PRODUCTIVITY GAINS, DESKILLING RISKS AND THE ROLE OF CONTINUOUS EMPLOYEE EDUCATION

Tamara Papic^{1*} 

¹Faculty of Technical Sciences, Singidunum University, Belgrade, Serbia, e-mail: tpapic@singidunum.ac.rs

Abstract: The pace of AI adoption in today's business world has sparked a debate about the dual nature of AI and its effect on human productivity. AI tools are great for work, but they also make employees overreliant, with weaker skills and less critical thinking. This paper examines the advantages of AI tools at the workplace level, together with the issues of cognitive offloading, deskilling and low learning motivation among employees. We conducted a systematic review of the literature on AI coding apps (GitHub Copilot, Cursor, etc.) and on corporate e-learning platforms (Coursera and Udemy) from the past several years. The review shows that these tools bring real productivity gains - in controlled experiments between 14% and 40% faster task completion - and that the gains are strongest for less experienced employees. At the same time, the same literature documents growing evidence of cognitive offloading and deskilling when AI assistance is used uncritically. The outcome of this research is that AI, combined with continuous education and the development of a learning culture (microlearning, MOOC-based upskilling and guided on-the-job learning), keeps employees motivated and more productive, while deskilling is reduced. The paper concludes that AI tools should not replace competence but should add to it. Management should build AI usage policies with clear ongoing education programs for employees. These findings matter to human resource managers, executives and policymakers who plan workforce development in the digital age.

Keywords: *artificial intelligence, workplace productivity, deskilling, continuous education, upskilling.*

Field: Social Sciences

1. INTRODUCTION

In the large field study by Brynjolfsson et al. (2025), a customer support agent who could not solve an issue no longer had to wait for a senior colleague to step in. The same agent received suggestions from a generative AI assistant, and resolved on average 14% more problems in one hour. The gain was more than a third for the least experienced agents. Numbers of this size are why AI assistants stopped being experiments a long time ago. They write reports, condense long email correspondence, generate code and quietly make small decisions that used to be the hands-on work of a person. Managers see the impact on quarterly dashboards. Dashboards do not show what is happening in between, especially with the knowledge of the people who use these tools.

The worry is not new. It is becoming harder to ignore. When people continuously let a system remember, reason and solve difficulties for them, these abilities are simply less exercised, and skills that are not exercised tend to fade. The mechanisms that drive the process are labeled cognitive offloading (Risko & Gilbert, 2016). Its more serious, longer-term form is now called deskilling (Subramanya et al., 2026). Why this price has to be paid at all is a different question. It should not be. The evidence from corporate training and workforce learning programs suggests that companies that treat AI adoption and employee upskilling as a single investment and not as two separate budget lines seem to have more productivity gains while minimizing skills loss (Al Naqbi et al., 2024; Rosendale & Wilkie, 2021).

RQ1: how big are the productivity gains of AI tools under working conditions? RQ2: what deskilling risks have been reported when workers rely on AI assistance? RQ3: can continuous employee education, which is deliberately implemented as a management tool, reduce these risks? Section 2 describes how the literature was taken and studied. Section 3 presents the empirical results. Section four discusses implications and offers an integrative model. Section five finishes with conclusions.

*Corresponding author: tpapic@singidunum.ac.rs



2. MATERIALS AND METHODS

We conducted the study as a systematic review of research literature and published research. We sought sources from scholarly databases and publisher platforms (Google Scholar, MDPI, Springer, Sage, Emerald, Taylor & Francis, IGI Global) and the main search period was from 2021-2026. We chose to address two older but foundational sources, Risko and Gilbert (2016), who gave cognitive offloading its current theoretical picture, and Acemoglu and Restrepo (2019), who covered automation in the context of the task-based approach.

We searched for four categories of research topics: AI and workplace productivity; AI coding assistants such as GitHub Copilot, Cursor, ChatGPT; deskilling, skill erosion and cognitive offloading; and corporate e-learning, MOOC-based upskilling and continuing employee training. A source was included only if it had been peer-reviewed or documented a field experiment, if it related to AI use at work or to employee education, and if it was written in English. The practice papers and opinion pieces without review were left out.

Thematic synthesis followed. The remaining sources were sorted into three categories (measured productivity effects, deskilling mechanisms and educational countermeasures) and the quantitative results of field experiments were extracted and compared across task types and experience levels. Why is there a review rather than actual data? AI tools move faster than most organizations can study, and single company measurements and time-series data can become very old in months; it is a synthesis of controlled experiments that supports managerial conclusions. One caveat should be noted at the start. Most of the available experiments are conducted by North American and Western European firms and the results do not come out in the same way for smaller labor markets, such as Southeast Europe, where adoption and training infrastructure are much different (Premović et al., 2025).

3. RESULTS

Productivity first. And the strongest evidence is collected through experiments rather than surveys. Noy and Zhang (2023) randomized 453 professionals and found that having a generative AI assistant did cut the time spent writing by 40%, while external graders rated output 18% better. Another important point to keep in mind here: these were bounded writing assignments, and the authors themselves warned against extrapolating to open-ended work. Brynjolfsson et al. (2025) followed more than 5,000 customer support agents and found an average 14% increase in resolved problems per hour. The figure that managers should not miss is below the average: novice and low-educated agents increased 34% in response, suggesting that the assistant enables experts to share some of the knowledge that is otherwise not available to the beginner. Cui et al. (2026) tracked roughly 4,800 software developers at Microsoft, Accenture and a Fortune 100 company who used GitHub Copilot, and they found 26% more completed tasks. Since this study followed real work over a long period of time, not a short laboratory session, it has more weight than most studies. Finally, the PRISMA-based review of Al Naqbi et al. (2024) covers 159 publications from different sectors, and finds the same direction of effects across contexts, though it also lists a number of unresolved risks


Table 1. Overview of key empirical studies on generative AI and workplace productivity

Study	Context / tool	Sample	Measured effect
Noy & Zhang (2023)	Professional writing, ChatGPT	453 professionals, randomized	40% less time per task; +18% output quality
Brynjolfsson et al. (2025)	Customer support, GenAI assistant	5,179 support agents	+14% issues/hour; +34% for novices
Cui et al. (2026)	Software development, GitHub Copilot	~4,800 developers, 3 large firms	+26% completed tasks (field experiment)
Al Naqbi et al. (2024)	PRISMA literature review, GenAI	159 publications, cross-sector	Consistent gains; persistent risks

Source: Compiled by the author based on Noy and Zhang (2023), Brynjolfsson et al. (2025), and Cui et al. (2026).

Table 1 and Figure 1 summarize our findings. The pattern is clear: the gains are real, but not shared equally. Junior employees benefit the most; experienced staff sometimes gain little and on tricky tasks even lose ground. An organization that introduces AI tools without changing its training plan might get the new people fast while doing nothing for its most experienced ones.

Figure 1. Measured productivity effects of generative AI tools in controlled studies



Source: Compiled by the author based on Noy and Zhang (2023), Brynjolfsson et al. (2025), and Cui et al. (2026).

The second cluster of results is not so comfortable. In cognitive offloading theory (Risko & Gilbert, 2016), cognitive skills that are regularly handed over to an external system are undermined by disuse. Subramanya et al. (2026) explore what they call a complementarity paradox: working with AI initially improves performance but overreliance over time erodes the ability to perform tasks unaided. Their argument is persuasive precisely because it identifies a mechanism, not a correlation - the decline is a predictable consequence of reduced cognitive engagement. Lovett (2026) adds an important qualification: deskilling and inhibited upskilling cluster in particular user profiles so that the risk, while real, is not uniform across a workforce. In software engineering, the birthplace of tools like GitHub Copilot and Cursor, Le (2026) reports that AI assistants change workflows and daily decisions in ways that enhance output

today but leave open questions about tomorrow - how developers will keep debugging and architectural reasoning.


The third cluster turns to education as an alternative. Corporate e-learning platforms (Coursera and Udemy are two of the best known) along with MOOC-based programs offer organizations a scalable way to reskill and upskill their employees (Rosendale & Wilkie, 2021). Alasmari (2024) reveals that employees actually use such programs depending on their perceived usefulness and organizational support; buying platform licenses should not be the same as building a learning culture. AI adoption in software teams follows the same trajectory from the other direction, with structured training acting as an efficiency multiplier only where the tools are used (Koivisto et al., 2026).

4. DISCUSSION

Putting the two literatures next to each other changes the question. It's not about whether AI tools are good or bad for a workforce. The productivity studies show quick, measurable output gains; the deskilling studies show that those gains can quietly consume the capabilities the output depends on. The two effects do not cancel out - they run on different clocks, which is why the damage is not really visible in quarterly reports. What connects them is learning. Employees who are trained, systematically, to understand, check and challenge what the AI gives them tend to keep and improve on what it has taught them and even to extend it. Employees who are passive with the same tools perform worse. One of the most important findings of this review that I want to emphasize is that the same tool has the opposite long-term impact depending on the educational environment around it.

Figure 2 shows the model we propose. AI adoption has two parallel effects - productivity gains on one hand; deskilling risks on the other. Continuous employee education slows the risk path and strengthens the gain path, turning a brief output spike into long-term growth. The practical result is plain but vital: decisions about buying AI tools and decisions about training budgets should happen in the same meeting. In our experience with organizations in the region, they are still almost always held separately.

Figure 2. Integrative model: AI adoption, deskilling risk and the moderating role of continuous education



Source: Author's own elaboration.

In software teams, assistants such as GitHub Copilot and Cursor speed up coding. Teams that teach developers to check generated code carefully can keep competence (Koivisto et al., 2026; Le, 2026). Outside IT, MOOC platforms such as Coursera and Udemy let companies offer microlearning, closing skill gaps for far less than traditional training costs (Rosendale & Wilkie, 2021). In every case the same lesson remains: AI output should be seen as learning material, not as a finished product.

All of this is consistent with the task-based view of automation (Acemoglu & Restrepo, 2019). Technology takes over some tasks and creates or reinstates others, and where the balance comes from will be dependent on complementary investments. So far, of all such investments it seems that workforce education is the one that managers are able to most directly control, and according to studies we have reviewed here, the one that is the least frequently covered.

5. CONCLUSIONS

We asked three things here: what AI tools can add to productivity, what they can take away from employee competence, and whether education can reconcile the two. The productivity evidence is unambiguous. Controlled studies show output increases of 14% to 40% in low-skill employees. The deskilling evidence is younger but also consistent across all studies: uncritical use of AI promotes cognitive offloading and gradually erodes the skills employees need when they do not have the tool. On the third question, the answer is actually very encouraging. Organizations that integrate learning into their AI deployment - not just bolt it on afterwards - will more likely keep their productivity gains going for a longer time.

The management advice is in this sense so simple: consider AI as an augmentation to human skills, have mandatory and measurable training for every deployment, and measure skill retention in the same manner as output metrics. In Southeast Europe, where training budgets are typically so low and adoption is so rapid, this sequence matters much more than in the markets where the reviewed experiments were performed. There are limits to the study that need to be noted. It is based on published experiments, mostly from writing, customer support, and software development. Longitudinal field studies of other sectors with the ability to measure skill retention in real time are the next step.

REFERENCES

- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3-30. <https://doi.org/10.1257/jep.33.2.3>
- Alasmari, T. (2024). MOOCs for reskilling and upskilling: A perspective of employees' acceptance. *International Journal of Multi Discipline Science (IJ-MDS)*, 7(2). <https://journal.stkip singkawang.ac.id/index.php/IJ-MDS/article/view/6260>
- Al Naqbi, H., Bahroun, Z., & Ahmed, V. (2024). Enhancing work productivity through generative artificial intelligence: A comprehensive literature review. *Sustainability*, 16(3), 1166. <https://doi.org/10.3390/su16031166>
- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889-942. <https://doi.org/10.3386/w31161>
- Cui, K. Z., Demirel, M., Jaffe, S., Musolff, L., Peng, S., & Salz, T. (2026). The effects of generative AI on high-skilled work: Evidence from three field experiments with software developers. *Management Science*. Advance online publication. <https://doi.org/10.1287/mnsc.2025.00535>
- Koivisto, H., Koivusaari, J., & Saarela, K. (2026). Adoption of AI-assisted code generation in software development: Insights from GitHub Copilot trainings. In *Proceedings of the Fifth International Conference on Innovations in Computing Research (ICR'26) (Lecture Notes in Networks and Systems*, pp. 180-194). Springer Nature Switzerland. https://doi.org/10.1007/978-3-032-32636-2_16
- Le, H. (2026). Impact of generative AI assistants on software engineers' skills, workflow efficiency, and decisions. In *AI impacts on deskilling and reskilling software engineers* (pp. 1-26). IGI Global Scientific Publishing. <https://doi.org/10.4018/979-8-3373-8450-4.ch001>
- Lovett, N. (2026). Individual differences in generative AI effectiveness: An empirical puzzle for HRD research. *New Horizons in Adult Education and Human Resource Development*. Advance online publication. <https://doi.org/10.1177/19394225261472792>
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187-192. <https://doi.org/10.1126/science.adh2586>
- Premović, J., Krmpot, V., & Premović, T. (2025). The implementation of artificial intelligence in the educational institutions: evidence from Serbia. *SCIENCE International Journal*, 4(4), 123-128. <https://doi.org/10.35120/sciencej0404123p>
- Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676-688. <https://doi.org/10.1016/j.tics.2016.07.002>
- Rosendale, J., & Wilkie, L. A. (2021). Scaling workforce development: Using MOOCs to reduce costs and narrow the skills gap. *Development and Learning in Organizations: An International Journal*, 35(2), 18-21. <https://doi.org/10.1108/DLO-11-2019-0258>
- Subramanya, S., Krishnan, S., & Azad, N. (2026). The complementarity paradox: Human-AI collaboration, overreliance, and skill decline in entrepreneurial contexts. *Information Systems Frontiers*. <https://doi.org/10.1007/s10796-026-10804-5>