

Generative AI and the Nature of Work

IN THIS BRIEF

- Artificial intelligence has been shown to improve human productivity, but could the technology also change the nature of work itself? To explore this question, a team of five researchers conducted one of the largest-ever natural experiments of Generative AI.
- Over the course of two years—July 2022 to July 2024—the researchers observed the coding work of more than 187,000 software developers using GitHub Copilot, a GenAI tool for software development.
- The researchers focused on open source in part because these developers typically work in decentralized settings. They're also often overburdened, spending more time than they'd like on project development and less time than they'd like on actual coding.
- The researchers found that top developers who received free access to GitHub Copilot during the observed period increased their coding tasks as a share of all tasks they performed. The developers also reduced their relative share of activities related to project management.
- In addition, GenAI was found to be most helpful for open source developers with lower levels of coding experience. During the observed period, lower-ability developers using GitHub Copilot increased their coding and decreased their project management more than their highly skilled counterparts
- The experiment shows that GenAI holds great promise for allowing open source developers to spend more time in the manner they prefer—writing code—while also ensuring their software's security, stability and usability. What's more, as distributed work becomes increasingly common, this effect is likely to generalize to other occupations and settings as well.

RESEARCH OVERVIEW

A mere handful of technological innovations—among them the printing press, internal combustion engine and general-purpose computer—have fundamentally changed the way people live and work. Given recent advances, artificial intelligence may join this elite category of technologies (Crafts, 2021; Goldfarb et al., 2023; Eloundou et al., 2024). AI's greatest economic impact could be improving productivity in knowledge-intensive industries (Manyika et al., 2018; Sachs, 2023).

However, research into AI's impact is still nascent. This is especially true for Generative AI, a subset of the technology built on large language models (LLMs). Early studies have shown that GenAI can make high-level impacts on productivity (Brynjolfsson et al., 2023; Dohmke et al., 2023; Noy & Zhang, 2023; Peng et al., 2023). Less clear, however, are the mechanisms driving these improvements.

One hint comes from prior research showing how technologies that streamline communication and decision-making processes can reduce the overhead of collaboration, freeing workers to focus on their own work in isolation (Faraj et al., 2011; Aral & Van Alstyne, 2011). GenAI takes that process a step further. With the technology, many of these collaborative costs are simply eliminated. Work that previously required communication among multiple people can now be done without any interaction at all.

To explore these and related issues, a team of five academic and business researchers—Manuel Hoffman, Sam Boysel, Frank Nagle, Sida Peng and Kevin Xu—designed and conducted a natural experiment.

They sought answers to questions including:

- What is the effect of AI technology on task allocation across specific kinds of core work and project management?
- When workers use AI, do they favor exploitation or exploration in task allocation? That is, are they more likely to increase their efforts on projects they are familiar with? Or are they more likely to branch out into projects that they've never worked on before?
- Is it true that AI helps lower-ability workers more than it helps higher-ability workers?

The researchers describe their experiment and its results in a recent working paper, [Generative AI and the Nature of Work](#).

THE EXPERIMENT

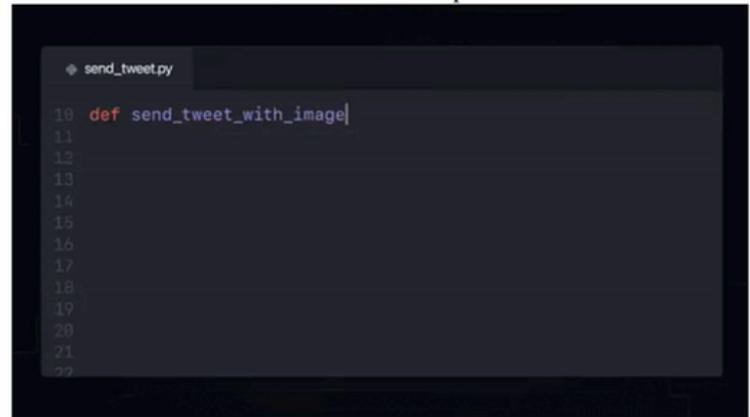
To conduct their natural experiment, the researchers first needed a setting with two shared characteristics: One, the work is done from distributed locations, and two, the work tasks of individuals can be observed in great detail. As the researchers discovered, both requirements could be met in open source software development.

The researchers also needed a before-and-after setting in which the effects of a new AI tool could be clearly observed. To meet this additional requirement, the researchers selected the June 2022 public release of GitHub Copilot, an AI software-development tool.

GitHub offered several clear benefits. As the world's largest hub for open source software developers, GitHub provides cloud-based services for both software development and version control. More to the point, GitHub has been designed for use by geographically dispersed teams. Also, GitHub documents all activities performed on its system. This allowed the researchers to observe in granular detail the work tasks completed by remote teams of software developers, making it an ideal setting for their natural experiment.

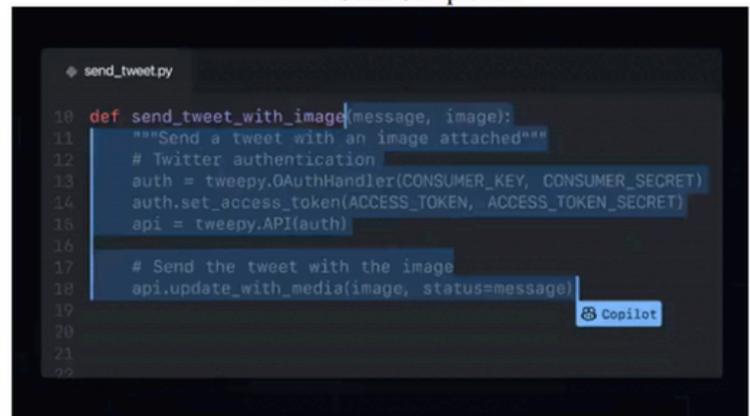
The researchers opted to study the impact of one specific GitHub tool: Copilot, a GenAI software-development tool developed jointly by GitHub, OpenAI and Microsoft. While GitHub Copilot is based on a predictive model similar to that used by ChatGPT, the tool differs in important ways from both ChatGPT and Microsoft's similarly named Copilot tool. Developers using GitHub Copilot can generate code snippets that are easily integrated into existing codebases (Fig. 1). It's a popular approach; in one recent survey, nine in 10 U.S.-based developers said they use an AI coding tool (Shani, 2023).

Panel A: User Input



```
send_tweet.py
10 def send_tweet_with_image|
11
12
13
14
15
16
17
18
19
20
21
22
```

Panel B: Code Completion



```
send_tweet.py
10 def send_tweet_with_image(message, image):
11     """Send a tweet with an image attached"""
12     # Twitter authentication
13     auth = tweepy.OAuthHandler(CONSUMER_KEY, CONSUMER_SECRET)
14     auth.set_access_token(Access_token, ACCESS_TOKEN_SECRET)
15     api = tweepy.API(auth)
16
17     # Send the tweet with the image
18     api.update_with_media(image, status=message)
19
20
21
22 Copilot
```

Figure 1: GitHub Copilot in action. First, the human developer wrote a function (Panel A). Then, based on this prompt, Copilot suggested the rest (Panel B).

Source: GitHub, 2022

The natural experiment consisted of a panel of 187,489 developers, and the researchers observed these developers on a weekly basis from July 2022 to July 2024. Over this two-year period, the researchers made literally millions of observations of the developers' weekly work.

The researchers focused specifically on developers deemed "top maintainers" by GitHub, which made them eligible for free access to the Copilot tool. (Most developers can use GitHub Copilot for free only during short trial periods. Thereafter, they must pay a monthly fee.)

The researchers further organized their observations along two main categories of essential developer work: coding and project management. Under coding they included the more technical processes of writing lines of software code.

Project management was the heading for most remaining activities, including assisting other developers with software issues, introducing new ideas to the developer community, and discussing long-term objectives (Fig. 2).

There are other tasks a developer can do that do not fall in one of these buckets; so just because coding increases as a share of all activities, project management doesn't necessarily need to decrease.

Coding	Project Management
Create Repository	Created Project Board
Fork Repository	Issue Assigned
Pull Request	Issue Closed
Push	Issue Comment Deleted
	Issue Closed
	Issue Labeled
	Issue Merged
	Issue Reopened
	Issue Review Requested
	Issue Subscribed
	Reviewed Pull Request

Figure 2: Classification of software developers' work activities. Each category is defined as the sum of its disaggregated, granular activities.

THE RESULTS

The researchers first established that GitHub's program for top developers increases Copilot usage for eligible users. Top developers used Copilot significantly more than other Copilot-adopting developers did, and more of them adopted Copilot.

Next, the researchers explored the causal impact of access to Copilot on patterns of distributed work. Overall, they found that GenAI induces developers to reallocate toward core work. Among the top developers observed, their coding work as a percentage of all activity increased by 5.4%, while their project management work as a percentage of all activity decreased by 10%. This also implies that developers with access to GenAI tools are less likely to seek help from other developers. Instead, they use the GenAI tool to address their problems or inquiries.

Due to the long-term nature of the natural experiment, the researchers could also examine whether the developers' use of Copilot changed over time. It turned out that the strongest effects took hold during the first year. Then, after some experimentation, the impacts were stable for approximately two years (the end of the study).

Another question answered by the natural experiment was whether GenAI induces developers to work more autonomously. Overall, the data showed that Copilot allowed developers to work by themselves more; therefore, they worked with others less. That was mainly a result of the GenAI tool allowing the developers to spend more time on their core (and more solitary) activity of coding.

Yet another question the researchers explored was whether GenAI encourages developers to branch out into experimental work that transcends their established projects. The short answer is, it does. The researchers found that on average, Copilot-eligible developers engaged with 15 more new projects than did their ineligible peers. These same developers also increased their exposure to new programming languages by nearly 22% relative to the baseline.

One final question was whether GenAI helps developers of all ability equally. To answer this, the researchers used measures that included project workload, contribution diversity, and popular interest from peers. They found that Copilot helped low-ability developers more than it did those of high ability.

CONCLUSIONS

This natural experiment has several broader implications. One is that managers may be aided by GenAI's ability to change the makeup of work. For many organizations, GenAI may bring about more streamlined production processes. Another is that GenAI technology has the potential to flatten organizational hierarchies. Yet another is that talented workers may use GenAI to refocus their work on both core processes and new, exploratory innovation.

The researchers also believe that GenAI can help software developers earn more money. A developer can use GenAI tools to gain exposure to new programming language, and the researchers say that new skill can increase a developer's earning potential by an estimated \$1,683 a year. That figure,

multiplied by the full set of 300,000 developers working on open source, suggest that Copilot could improve their combined annual income by as much as \$468 million.

The researchers concede this is a back-of-the-envelope calculation. Yet they believe that given GenAI-powered productivity improvements and other experimentation, the true value could be even higher.

REPORT

Read the [full working paper](#)

ABOUT THE AUTHORS

Manuel Hoffman is an Assistant Professor at the University of California, Irvine, and a former Postdoctoral Fellow at Harvard Business School.

Sam Boysel is a Data Scientist at The Linux Foundation and a former Postdoctoral Fellow at Harvard University's Laboratory for Innovation Science.

Frank Nagle is a Research Scientist with the MIT Initiative on the Digital Economy (IDE); Advising Chief Economist for The Linux Foundation; and a former Assistant Professor at Harvard Business School.

Sida Peng is a Research Economist in the Office of the Chief Economist at Microsoft Research.

Kevin Xu is a Staff Software Engineer at GitHub Inc.

REFERENCES

Aral, S., and Van Alstyne, M. (2011). [The diversity-bandwidth trade-off](#). *American Journal of Sociology*, vol. 117, no. 1, pp. 90-171.

Brynjolfsson, E., et al. (2023). [Generative AI at work](#). National Bureau of Economic Research, working paper 31161.

Crafts, N. (2021). [Artificial intelligence as a general-purpose technology: an historical perspective](#). *Oxford Review of Economic Policy*, vol. 37, issue 3, pp. 521-536.

Dohmke, T., et al. (2023). [Sea change in software development: economic and productivity analysis of the AI-powered developer lifecycle](#). arXiv:2306.15033.

Eloundou, T., et al. (2024). [GPTs are GPTs: Labor market impact potential of LLMs](#). *Science*, vol. 384, issue 6702, pp. 1306-1308.

Faraj, S., et al. (2011). [Knowledge collaboration in online communities](#). *Organization Science*, vol. 22, no. 5, pp. 1224-1239.

Goldfarb, A., et al. (2023). [Could machine learning be a general purpose technology? A comparison of emerging technologies using data from online job postings](#). *Research Policy*, vol. 52, issue 1, article 104653.

Manyika, J., et al. (2018). [The promise and challenge of the age of artificial intelligence](#). McKinsey Global Institute briefing note.

Noy, S., and Zhang, W. (2023). [Experimental evidence on the productivity effects of generative artificial intelligence](#). *Science*, vol. 381, issue 6654, pp. 187-192.

Peng, S., et al. (2023). [The impact of AI on developer productivity: Evidence from GitHub Copilot](#). arXiv:2302.06590.

Sachs, G. (2023). [Generative AI could raise global GDP by 7%](#). Goldman Sachs article.

Shani, I. (2023). [Survey reveals AI's impact on the developer experience](#). GitHub survey.



MIT Initiative on the Digital Economy

MIT Sloan School of Management
245 First St, Room E94-1521
Cambridge, MA 02142-1347

ide.mit.edu

Our Mission: The MIT Initiative on the Digital Economy (IDE) is shaping a brighter digital future. We conduct groundbreaking research on the promise—and peril—of new digital technologies including generative artificial intelligence (GenAI), quantum computing, data analytics, and distributed marketplaces. We also investigate the rise of fake news and misinformation and the development of a digital culture. Through research and the convening of leaders from academia, industry, and government, the IDE provides critical, actionable insight for people, businesses, and government to understand and benefit from new technologies and how they're rapidly changing the ways we live, work, and communicate.

Contact Us: David Verrill, Executive Director,
MIT Initiative on the Digital Economy
617-452-3216
dverrill@mit.edu

Become a Sponsor: The generous support of individuals, foundations, and corporations help to fuel cutting-edge research by MIT faculty and graduate students. It also enables new faculty hiring, curriculum development, events, and fellowships.

Additional Contact: Albert Scerbo
Associate Director
MIT Initiative on the Digital Economy
267-980-2616
ascerbo@mit.edu

[View all our sponsors](#)

Connect with us:

