

Evolving with AI: A Longitudinal Analysis of Developer Logs

Agnia Sergeyuk*
JetBrains Research
Belgrade, Serbia
agnia.sergeyuk@jetbrains.com

Anastasiia Serova
JetBrains
Amsterdam, Netherlands
anastasiia.serova@jetbrains.com

Eric Huang*
University of California, Irvine
Irvine, United States
jianch2@uci.edu

Yaroslav Golubev
JetBrains Research
Belgrade, Serbia
yaroslav.golubev@jetbrains.com

Dariia Karaeva
JetBrains
Paphos, Cyprus
dariia.karaeva@jetbrains.com

Iftekhhar Ahmed
University of California, Irvine
Irvine, United States
iftekha@uci.edu

Abstract

AI-powered coding assistants are rapidly becoming fixtures in professional IDEs, yet their sustained influence on everyday development remains poorly understood. Prior research has focused on short-term use or self-reported perceptions, leaving open questions about how sustained AI use reshapes actual daily coding practices in the long term.

We address this gap with a mixed-method study of AI adoption in IDEs, combining longitudinal two-year fine-grained telemetry from 800 developers with a survey of 62 professionals. We analyze five dimensions of workflow change: productivity, code quality, code editing, code reuse, and context switching. Telemetry reveals that AI users produce substantially more code but also delete significantly more. Meanwhile, survey respondents report productivity gains and perceive minimal changes in other dimensions.

Our results offer empirical insights into the silent restructuring of software workflows and provide implications for designing future AI-augmented tooling.

CCS Concepts

- **Human-centered computing** → **Empirical studies in HCI**;
- **Software and its engineering** → *Software maintenance tools*; *Integrated and visual development environments*.

Keywords

Developer behavior, AI assistants, IDE, Productivity, Code quality, Code editing, Code reuse, Context switching

ACM Reference Format:

Agnia Sergeyuk, Eric Huang, Dariia Karaeva, Anastasiia Serova, Yaroslav Golubev, and Iftekhhar Ahmed. 2026. Evolving with AI: A Longitudinal Analysis of Developer Logs. In *2026 IEEE/ACM 48th International Conference on Software Engineering (ICSE '26)*, April 12–18, 2026, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3744916.3787811>

*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution 4.0 International License. *ICSE '26, Rio de Janeiro, Brazil*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2025-3/2026/04
<https://doi.org/10.1145/3744916.3787811>

1 Introduction

Recent advances in Large Language Models (LLMs) have introduced a profound shift in how software is developed [11, 23, 40]. AI-powered coding assistants such as GitHub Copilot, JetBrains AI Assistant, and Visual Studio IntelliCode are now embedded in the daily workflows of many professional developers [18, 38]. These tools are designed to automate routine coding tasks, generate non-trivial code suggestions, and support activities like documentation and test generation [21, 34]. However, while these tools promise productivity and efficiency gains, their long-term effects on how developers work — how they write, edit, reuse, and navigate code — remain poorly understood.

Understanding these shifts is not only timely but essential. As AI tools are rapidly adopted, they have the potential to fundamentally restructure developer workflows, with implications for productivity, software quality, team dynamics, and future tooling. Most existing studies on developer-AI interaction rely on short-term experiments [16, 27, 42], small-scale observations [8, 11], or surveys of developer perceptions [21, 41]. While valuable, these approaches provide a limited view of the evolving developer experience. Lab studies may not generalize to the complexity and scale of real-world software development, and self-reported data can differ notably from actual behavior [4, 12]. Importantly, most prior work offers only a static view, missing the longer-term evolution of developer practices as they adapt to and co-evolve with AI tools.

We address this gap through a comprehensive, mixed-methods analysis of developer behavior, conducted in collaboration with JetBrains¹, a major software vendor specializing in the creation of intelligent development tools. Our study combines longitudinal, large-scale telemetry data with qualitative insights. We analyze 151,904,543 fine-grained, anonymized IDE interaction logs collected over two years from 800 software developers, including 400 *AI users* and 400 *AI non-users*. To complement this behavioral data with self-reported experiences and perspectives, we conducted a survey with 62 professional developers and follow-up interviews with five participants. This approach examines developer workflows through two complementary lenses: what developers do (as observed in logs), and what they perceive (as captured through surveys and interviews), offering a detailed view of AI's real-world impact on software development practices.

The research questions guiding this study are as follows:

How did the introduction of AI tools influence the...

¹JetBrains: <https://www.jetbrains.com/>

RQ1: ...*productivity of developers?*

RQ2: ...*code quality?*

RQ3: ...*practices of code editing?*

RQ4: ...*practices of code reuse?*

RQ5: ...*practices of context switching?*

Together, these five dimensions provide a structured way to examine how AI assistance influences not only the outcomes of software development but also the underlying practices and processes. By addressing these five questions, we aim to develop both a fine-grained and holistic understanding of how developer workflows are evolving in response to sustained AI tool usage.

Our results offer implications for the future of software engineering practice and tooling.

- **Empirical characterization of evolving AI-assisted workflows.** Through examination of logs, we reveal that developers using AI tools write and modify significantly more code, rely more frequently on external sources (e.g., code copied from the web), and switch contexts more often. At the same time, survey respondents perceive minimal changes outside of productivity.
- **Reframing of AI's impact on effort and attention.** Our findings show that AI does not uniformly reduce developer effort; instead, it redistributes effort across more fragmented, reactive, and cognitively demanding workflows.

Together, these findings deepen our understanding of how AI tools reshape software development in practice. They highlight the limitations of self-reported data, reveal structural changes in workflow that are not always visible to users themselves, and point to new design priorities for AI-assisted environments. Rather than focusing solely on task automation, future tools should also support continuity, reduce cognitive overhead, and enhance developer awareness across fragmented tasks.

The remainder of the paper is organized as follows. Section 2 reviews relevant literature and motivates the five workflow dimensions studied. Section 3 describes our mixed-methods approach, including data collection and analysis procedures. Section 4 presents our findings by research questions. Section 5 synthesizes the results and their implications. Section 6 discusses threats to validity, and Section 7 concludes the paper. The supplementary materials for the paper are available online [35].

2 Related Work

2.1 Developers' Workflows in IDEs

Understanding how developers spend their time in IDEs is a focus of interest for both software engineering research and industry. Prior work has leveraged telemetry to uncover behavioral patterns that reflect key aspects of developers' workflow. For instance, Minelli et al. conducted an in-depth analysis based on data mined from developers' interactions within the PHARO IDE [22]. From this, they discovered that developers spend the majority of their time (70%) on program comprehension activities, which include reading, navigating, and reviewing source code. Similarly, Amann et al. [2] analyzed usage traces from Visual Studio, showing that developers interact with the IDE for only 25% of their time, with frequent brief interruptions punctuating longer work sessions.

To move beyond low-level metrics, researchers have employed sequence modeling and clustering techniques to abstract developers' behaviors. Damevski et al. [9] used Hidden Markov Models to categorize actions related to debugging into latent task states. Astromskis et al. [3] analyzed over 1,000 hours of IDE interaction to define six high-level categories — code-related activities, executing the system under development, using utilities, using external programs, communicating with other developers, and looking for online help. Collectively, these studies demonstrate how telemetry enables scalable, in-situ observation of developer behavior.

While prior work has provided foundational insights into how developers engage with IDEs, it primarily captures behavior in pre-AI contexts. These studies offer valuable baselines for understanding time allocation, navigation patterns, and debugging workflows, but they do not account for the widespread adoption of AI-powered coding assistants in modern development environments. As such, they leave open critical questions about how AI alters not just developer productivity, but also the structure, cadence, and cognitive demands of everyday development work. Our study addresses this gap by leveraging fine-grained IDE telemetry logs collected over a two-year period from professional developers actively using an AI assistant. Unlike earlier studies focused on static or low-level patterns, we combine these logs with developer survey responses to investigate how behaviors and perceptions evolve over time.

2.2 AI and Developers' Workflows in IDEs

The recent rise of LLMs has introduced new and potentially transformative dynamics into software development workflows. AI-based coding assistants, such as GitHub Copilot [14] and JetBrains AI Assistant [17] offer promises of increased efficiency and reduced manual effort. However, their actual impact on day-to-day development behavior remains unclear. Notably, most prior work has relied on short-term controlled experiments [16, 27, 42] or self-reported perceptions [21, 41], with limited use of longitudinal, telemetry-based evidence. To systematically investigate how AI affects real-world development practices, we focus on five workflow dimensions that are both grounded in prior research and observable through fine-grained telemetry from an enterprise-scale IDE: **productivity**, **code quality**, **code editing**, **code reuse**, and **context switching**.

Productivity. Since the emergence of LLM-based coding tools, numerous studies have explored their impact on developer productivity [4, 7, 8, 16, 29, 41, 42, 44, 46]. However, findings have been mixed, and there remains no clear consensus on whether such tools consistently yield measurable productivity gains in real-world development contexts. Some studies report promising results. For example, Imai [16] found that pair programming with GitHub Copilot led to increased code output — as measured by lines added — suggesting that AI assistance can accelerate the act of writing and modifying code. Similarly, Peng et al. [27] demonstrated that Copilot significantly reduced task completion time, yielding over a 50% improvement in productivity. However, other work urges caution. Vaithilingam et al. [42] observed that Copilot did not consistently reduce task time or improve task success in realistic programming tasks. Ziegler et al. [46] echoed this concern, highlighting a common disconnect between perceived productivity gains and actual improvements in development outcomes. In addition, Becker et al.

[4] found that, contrary to developers' own estimates, the use of AI assistants actually increases task completion time by 19%.

Code Quality and Code Editing. Prior research has examined how AI assistants are reshaping the programming process, revealing meaningful shifts in developer workflows. Given the imperfect nature of AI-generated code, developers now devote substantial effort to verifying, refining, and integrating AI suggestions. Mozannar et al. [23] report that developers spend over 50% of their time evaluating and editing AI-generated output. Complementary findings from Sahoo et al. [30] show that of the AI-suggested code that is initially accepted, 18.16% is later deleted and 6.62% is heavily rewritten — highlighting the costs of validation and rework.

Developers' attitudes toward AI assistance also vary. Some actively avoid using such tools due to concerns that generated code may not meet functional or non-functional requirements [21]. In contrast, others exhibit over-reliance, placing undue trust in AI output and showing reduced willingness to edit or critically assess it [28, 31]. This variation raises concerns about both under- and over-utilization of AI in the development process.

The impact of AI-generated code on code quality remains similarly inconclusive. Some studies document performance-related inefficiencies introduced by LLMs, such as deeply nested loops and redundant variable declarations [20]. However, other findings suggest that these issues may not substantially degrade the overall code quality. For example, Song et al. [37] report no significant change in quality despite observable productivity improvements. Weisz et al. [44] further argue that even when AI suggestions are suboptimal, they can function as effective scaffolding, shifting the task from code generation to code review.

Code Reuse. Code reuse — the practice of leveraging existing code, libraries, APIs, or prior implementations rather than writing functionality from scratch — has long been a cornerstone of efficient software development. However, the advent of AI-powered coding tools is reshaping reuse practices in subtle but significant ways. These tools frequently assist developers by generating boilerplate code and suggesting commonly used patterns or snippets derived from training data [5]. Recent empirical studies suggest that AI tools may both enable and displace traditional reuse strategies. Liang et al. [21] found that developers often rely on AI assistants for generating repetitive structures and simple logic, potentially reducing their dependence on existing libraries or shared components. At the ecosystem level, del Rio-Chanona et al. [10] reported a notable drop in StackOverflow activity — from approximately 60,000 to 40,000 new posts per month — within six months of ChatGPT release, suggesting a shift away from traditional forms of external knowledge reuse. Moreover, AI-generated code may unintentionally “reinvent the wheel” by replicating functionality already present in the codebase or available through standard libraries [36]. This raises concerns about the efficiency and sustainability of AI-facilitated reuse, particularly when such tools obscure opportunities to leverage mature, well-tested components.

Context Switching. One of the key promises of integrated AI coding assistants is their potential to reduce context switching by minimizing the need to leave the IDE to search for external solutions. By surfacing relevant suggestions and code completions

inline, these tools aim to make development workflows more seamless and less interruptive [26, 36]. However, emerging evidence suggests that new forms of disruption may offset these benefits. For example, Weber et al. [43] report that interacting with AI interfaces can introduce additional cognitive load and increase task fragmentation, as developers must frequently shift between writing code, interpreting suggestions, and managing communication with the AI system. These findings raise questions about whether AI assistance genuinely reduces context switching or replaces one form of interruption with another.

Summary. Despite the growing number of studies on AI in software engineering, the overall picture remains inconclusive. While some report perceived gains in productivity or faster code authoring, others highlight increased verification overhead or decreased code quality, resulting in a decrease in productivity. These mixed findings often stem from differences in study design, developer expertise, and task complexity, resulting in a fragmented and sometimes contradictory understanding of how AI tools reshape real-world development workflows.

To address this gap, we present a study of how AI tools affect developer workflows over a two-year period. Using anonymized IDE telemetry data from professional developers, combined with a survey capturing developer perceptions, we offer a comprehensive view of this transformation. Our analysis is grounded in five theoretically informed and operationally measurable dimensions of workflow. By complementing objective behavioral data with self-reported experiences, our study provides an integrated and empirically grounded perspective on how AI tools are reshaping software development practice over time.

3 Methodology

To study the influence of AI on developer workflows from different angles, we formulated research questions that focus on **how AI tools influence productivity, code quality, as well as the practices of code editing, code reuse, and context switching**. We performed a mixed-method study, consisting of two stages:

- **survey** of developers and the **post-survey interviews** that provided contextual insights;
- **log analysis** of actual developer behavior in the IDE.

This mixed-method approach examines AI adoption through two complementary but distinct lenses. Surveys capture retrospective, holistic evaluations of workflow change. Logs capture moment-to-moment interactions at the event level, aggregated monthly across thousands of actions. These sources measure different constructs and address the research questions from different temporal and granular perspectives. Our goal is to understand what developers actually do in their IDEs and what they perceive about those activities. Below, we describe each of the stages in detail.

3.1 Survey and Post-Survey Interviews

Firstly, to gather the self-reported perceptions from practitioners, we carried out an online survey aimed at our stated research questions. The study was conducted in line with our institution's ethical standards, adhering to the values and guidelines outlined in the ICC/ESOMAR International Code [1].

3.1.1 Data collection. Questions.

- Our survey consists of:
- Five **demographic questions**.
 - Two 5-point Likert-scale questions about the **overall experience** with AI tools for coding and the **reliance** on them.
 - Seven 5-point Likert-scale questions regarding **perception of their workflow** after they started using AI. Specifically, the questions focus on the changes to the five aspects of our study (from *Significantly decreased* to *Significantly increased*):
 - For **productivity**, we asked directly about **overall productivity** and also about **time spent coding (RQ1)**.
 - For **code quality**, we asked directly about the **quality of code** and also about **code readability (RQ2)**.
 - For **code editing**, we asked about the **frequency of editing or modifying own code (RQ3)**.
 - For **code reuse**, we asked about the **frequency of use of code from external sources**, e.g., libraries, online examples, or AI-suggested code (**RQ4**).
 - For **context switching**, we asked directly about the **frequency of context switching** — switching between different tasks or thought processes (**RQ5**).
 - An open-ended question where respondents were allowed to write a **specific example** of how AI tools for coding have impacted their workflow.
 - A question asking respondents if they want to participate in a short **post-survey interview**.

To ensure question clarity and unambiguity, we conducted two rounds of expert review with the institutional copyediting team and five test runs with colleague developers from JetBrains to validate survey question comprehension. The full questionnaire can be found in supplementary materials [35].

Participants. Recruitment was conducted via a survey link sent to a subset of developers who had previously consented to be contacted for research purposes as part of an internal participant panel maintained by JetBrains. This subset was selected based on prior survey responses in which they self-identified as users of AI-assisted development tools. These prior responses were collected as part of separate user studies conducted by JetBrains, with explicit consent for data retention and recontact. In line with internal research governance policies, no participants were contacted more than once within a 30-day period. A total of 1,231 invitation emails were sent to eligible participants, yielding 76 link clicks and ultimately 67 completed survey responses. As a thank you, participants were given the opportunity to enter a draw for USD 50 Amazon eGift Cards or an equivalent-value JetBrains product pack.

Overall, we received 67 complete survey responses. After excluding participants who indicated that they had not used AI assistance for coding, 62 responses remained for analysis. Of them, 56 identified as *Developers*, 16 as *Team Leads*, 10 as *Architects*, and 10 as *DevOps Engineers*. Participants could select up to three roles that described them. In terms of professional coding experience, 22 participants had *16 or more years*, 14 had *11–15 years*, 17 had *6–10 years*, 8 had *3–5 years*, and 1 had *1–2 years*.

Regarding AI assistance tools, the three most commonly used were *ChatGPT* ($n = 42$), *GitHub Copilot* ($n = 29$), and *JetBrains AI Assistant* ($n = 23$). Participants could select multiple tools. The distribution of AI tool usage duration was as follows: 28 participants

had used AI assistance for *more than one year*, 12 for *7–12 months*, 12 for *4–6 months*, 6 for *2–3 months*, and 4 for *less than one month* (*1–4 weeks*).

Post-survey interviews. In our survey, we asked participants for their consent to contact them for a post-survey interview to get more context on their answers. We selected five interview participants from survey respondents (P7, P15, P33, P39, and P62) who indicated interest in follow-up interviews. We aimed to capture diverse perspectives to ensure representative insights. We diversified across experience levels (3–5 years to 16+ years), geographic locations (Serbia, Chile, UK, Czech Republic, Argentina), AI tool satisfaction levels (from Very Dissatisfied to Very Satisfied), and usage patterns (regular users, occasional users, and those who stopped using AI tools). The five selected participants also represented different professional roles (Developer, DevOps Engineer, Architect, CIO), and AI tool usage durations (from weeks to over a year). This sampling approach captured varying experiences with AI adoption, including both positive and negative perspectives.

These participants were additionally rewarded with USD 50 Amazon eGift Cards or an equivalent-value JetBrains product pack upon their participation in the short (30 min), semi-structured interviews. The interviews focused on why participants reported specific perceptions and their reasoning behind them. We had JetBrains’s copyeditors review the interview script and performed a pilot session to refine question formulations and ensure clarity of the semi-structured interview approach. All interviews were conducted in English via Google Meet.² The full interview script can be found in supplementary materials [35].

We selected interview participants from survey respondents who indicated interest in follow-up interviews. We aimed to capture diverse perspectives to ensure representative insights. We diversified across experience levels (3–5 years to 16+ years), geographic locations (Serbia, Chile, UK, Czech Republic, Argentina), AI tool satisfaction levels (from Very Dissatisfied to Very Satisfied), and usage patterns (regular users, occasional users, and those who stopped using AI tools). The five selected participants represented different professional roles (Developer, DevOps Engineer, Architect, CIO), and AI tool usage durations (from weeks to over a year). This sampling approach captured varying experiences with AI adoption, including both positive and negative perspectives.

3.1.2 Data analysis. To characterize users’ perceptions of changes in their workflows, we calculated descriptive statistics for the questions regarding productivity, code quality, code editing, code reuse, and context switching.

For the answers to the open-ended question, we familiarized ourselves with their content and, based on it, invited participants for further short interviews. The interviews were transcribed into text using Google Meet’s built-in function and then validated by the authors. Similar to the answers to open-ended questions in the survey, the authors of the paper familiarized themselves with the transcripts and selected quotes that provided deeper context for the answers to the survey questions and overall findings.

²Google Meet: <https://meet.google.com/landing>

3.2 Log Analysis

To examine how developer behavior evolves over time, we carried out an analysis of IDE usage logs.

3.2.1 Data collection. The data about developer activities in the IDEs was provided by JetBrains in an anonymized form and consisted of the logs of specific actions (see details below). Each log consisted of a unique device ID (corresponding to a user), the name of the event, its timestamp, and various metadata.

Sample. The data consisted of usage logs from several IDEs, namely IntelliJ IDEA, PyCharm, PhpStorm, and WebStorm. Importantly, the studied IDEs covered different programming languages and ecosystems, representing a varied population of IDE users, not beholden to one particular language.

We gathered our samples from devices, for which the following two conditions were true:

- The device had some IDE activity in October of 2022.
- The device had some IDE activity in October of 2024.

The starting point of October 2022 was chosen because the first version of ChatGPT was released in November of 2022 [25], an event that enabled researchers and tool builders to start integrating LLMs into developers' tooling. For all the devices, we have the data for all the two years between these dates, allowing us to carry out a longitudinal analysis of changes that occurred over this time month-by-month.

We further draw two equal samples from this data:

- **AI users:** 400 devices that consistently interacted with JetBrains AI Assistant³ every month, at least once per month, from April 2024 to October 2024.
- **AI non-users:** 400 devices that never interacted with JetBrains AI Assistant during the studied two years.

The condition of using the in-IDE JetBrains AI Assistant since at least April 2024 allows us to ensure that *AI users* are people who have eventually adopted AI assistance into their development workflow. This window coincides with the period when AI assistants became widely available and stable in professional development environments, ensuring our *AI users* sample represents authentic adopters, who have integrated AI into their regular workflow rather than occasional experimenters. This allows us not only to look at logs but also compare AI adopters with other users, seeing how the rate of change for the studied aspects differed between the groups.

Metrics and actions. IDE telemetry logs are inherently complex and noisy, making it challenging to extract reliable insights without careful abstraction. To obtain robust and interpretable results, we focus on general, well-defined user actions that are consistently captured across sessions. Accordingly, we operationalize each of the five workflow aspects — productivity, code quality, code editing, code reuse, and context switching — in terms of measurable in-IDE behaviors recorded via telemetry.

While the generality of these actions means we cannot directly attribute observed changes solely to AI features, it is reasonable to assume that the introduction and increasing use of AI assistants

would influence these behaviors. Our goal is not to isolate causal effects, but rather to detect patterns of change in developer workflows over time, as AI tools become more integrated into everyday development practice. Moreover, having two samples of *AI users* and *AI non-users* allows us to see whether they statistically significantly differed in these general metrics.

Here are the actions and metrics we collected for each aspect:

- As a proxy for **productivity**, we counted the **number of typed characters in the editor (RQ1)**, reflecting the granularity of our available telemetry data. While code output has been used as an indicator of developer activity [13, 24], we acknowledge that it captures only one specific facet of productivity — namely, code authoring — and does not account for other important activities such as debugging, design, or collaboration. Nonetheless, it serves as a consistent and interpretable measure for analyzing changes in active code-writing behavior over time.
- As a proxy for **code quality**, we chose the **number of debugging instances (RQ2)**. Specifically, we collected all cases when developers started debugging (either by a hotkey or through contextual IDE actions). An increase in debugging activity may signal a rise in code defects or uncertainty, suggesting that developers are encountering more issues that require investigation.
- As a proxy for **code editing**, we chose the **number of deletions**, including delete keystrokes, backspaces, undos, etc. (**RQ3**). This aggregate measure reflects the frequency with which developers modify or iterate on existing code, providing insight into the intensity of editing activity. While it does not distinguish between types of edits, it serves as a general indicator of the effort devoted to refining code.
- As a proxy for **code reuse**, we chose the **instances of external content pasted into the IDE (RQ4)**. While paste actions can have diverse motivations and sources, the recent advancements in AI assistants, particularly chat-based interfaces, have increased copying from external suggestions [15]. For this RQ, we measure external code integration behavior rather than source-specific attribution. Therefore, we focused on paste actions that were not preceded by a copy action. This suggests that the content was not copied from within the codebase in the IDE but originated from external sources. This operationalization allows us to approximate the reuse of external code or AI-generated snippets.
- As a proxy for **context switching**, we chose the **number of IDE window activations (RQ5)**. This action refers to cases when the IDE was inactive (meaning the user was focused on another window, e.g., a browser), and then they switched back to the IDE. This metric captures how frequently developers shift their attention between the programming environment and external sources of information, providing insight into workflow fragmentation and interruptions.

Analyzing the dynamics of these actions enables a systematic observation of workflow evolution in *AI users* and *AI non-users*. Overall, our dataset comprised 151,904,543 logged events performed by the studied 800 users.

³JetBrains AI Assistant: <https://www.jetbrains.com/ai-assistant/>. This product is a general-purpose in-IDE AI assistant that integrates code completion, AI chat, code generation, and documentation features, capturing a diverse range of interaction patterns with AI inside IDEs.

3.2.2 Data processing. We collected longitudinal IDE telemetry data from October 2022 to October 2024 across all devices and conducted an analysis to understand how user behavior evolves over time. The raw data consisted of timestamped records of individual user actions. To reveal broader trends from this granular data, we aggregated actions on a monthly basis for each device.

Specifically, for each device and month, we computed the total count of occurrences for each tracked action. For example, we could determine that: *User A performed ‘Action X’ 15 times in January, 2023, and 22 times in February, 2023.* This process resulted in a clear, high-level dataset of monthly counts for every action, which is ideal for tracking activity and identifying meaningful behavioral trends.

No extensive filtering was applied to the data. In the cases when a particular user did not perform a particular action in the given month, we added it with a value of 0 to the list described above. This step ensured consistency across devices and time windows, preserving the full sequence for analysis.

3.2.3 Data analysis. Having obtained the final aggregated data, we moved on to analyzing it.

Normality and homogeneity checks. To inform the choice of statistical models, we tested for normality and homogeneity of variances. The Kolmogorov–Smirnov test [33] indicated that the distribution of action counts per device is not normal for all metrics. Bartlett’s test [32] showed that variances were not equal across groups and time points for most metrics. These results helped us choose the correct statistical tests.

Trend analysis. Having the aggregated monthly data for each metric, we carried out trend analysis on the longitudinal range from October 2022 to September 2024 (to cover exactly two years, 24 months). Given the non-normal distribution and heteroscedasticity of our data, we performed Mixed Linear Model Regression [39] for trend analysis. This approach is well-suited for data with repeated measures, where multiple observations are nested within subjects (in our case, devices). It allowed us to estimate changes over time, while accounting for inter-device variability. We fit a separate mixed-effects model for each of the five action types using the following formula:

$$count_{action} \sim group \times n_{month} + (1|id_{device})$$

Here, *group* indicates whether the device belongs to an *AI user* or an *AI non-user*, n_{month} is the numeric index of the month, and $(1|id_{device})$ specifies a random intercept for each device. This model enabled us to assess:

- Baseline differences between groups (*AI users* vs *AI non-users*).
- Overall temporal trends (e.g., whether the frequency of an action increases or decreases over time).
- Interaction effects, i.e., whether the trends differ significantly between groups.

Statistical significance was assessed using p-values, with a threshold of $p < 0.05$. You can find the full outputs of statistical testing, which highlight that our results are statistically significant, in supplementary materials [35].

Overall, this analysis enables a comparative examination of *AI users* and *AI non-users* with respect to how their coding practices

evolve over time. We examine these telemetry patterns alongside survey responses to provide complementary views of AI adoption expressed in practitioners’ perceptions of AI’s impact and the actual changes observed in their development workflows.

4 Results

In this section, we go over each of the five key aspects (**RQs 1–5**), drawing on the data from all the sources: survey, interviews, and log analysis.

Before focusing on the aspects, we also asked survey participants about their overall satisfaction with AI coding tools and their changing levels of reliance over time. The responses were largely positive. In terms of satisfaction, 71% of participants reported being satisfied or very satisfied, while only 11.3% reported dissatisfaction. Similarly, when asked whether their reliance on AI tools had changed since first adopting them, 61.3% indicated that their reliance had slightly or significantly increased, whereas only 14.5% reported a decrease. These results suggest that developers generally view AI tools favorably. However, such aggregate perceptions may obscure important nuances in how AI affects specific aspects of development work. To unpack these subtleties, we now examine perceptions around individual coding practices and juxtapose them with objective behavioral data from telemetry logs to answer the RQs. We present the findings in the following subsections.

4.1 RQ1: Productivity

The results for **productivity** are shown in Figure 1. Below, we first present findings from the survey responses, followed by insights derived from the telemetry log analysis.

Survey and interviews (Figure 1a). For 82.3% of respondents, the introduction of AI coding tools slightly or significantly increased their productivity, while two respondents said that it slightly or significantly decreased. Additionally, we asked about the influence of AI coding tools on time spent coding. Here, the results corroborate the same narrative: 56.5% of respondents said that their coding time decreased, and 14.5% indicated that it increased.

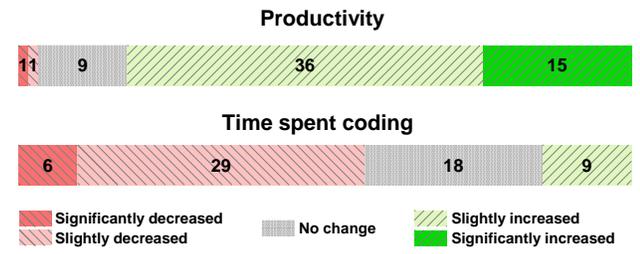
Judging these answers together, we can see the overall positive relation of participants towards productivity with AI-based tools. P7, a developer with 3-5 years of experience who regularly uses AI, provided an example of this:

“When I get stuck on naming or documentation, I immediately turn to AI, and it really helps.”

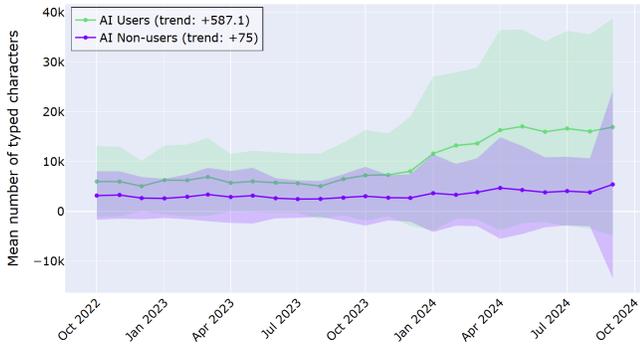
P33, a developer with more than 10 years of professional experience who churned from AI usage, on the other hand, provides a different testimony:

“I see a lot of time wasted trying to craft the right prompt to get the desired solution.”

Logs (Figure 1b). Telemetry data reveal that the number of typed characters used as a proxy for productivity increased over time for both *AI users* and *AI non-users*, with both trends being statistically significant. However, the rate of increase differs distinctly between the two groups. According to our mixed-effects linear model, *AI non-users* showed an average increase of 75 characters per month, whereas *AI users* exhibited a much steeper rise of 587 characters per month. This substantial difference suggests that AI adoption



(a) Responses to the survey questions whether *Productivity* and *Time spent coding* decreased or increased due to using AI tools for coding.



(b) The mean *Number of typed characters* per device in the IDE logs, for different months. *AI Users* show a statistically significant increasing trend of +587 characters a month, *AI Non-users* — +75 characters. Shaded regions represent ± 1 standard deviation from the monthly mean.

Figure 1: Results for RQ1: Productivity.

is associated with a greater acceleration in active code authoring over time.

Summary of RQ1. Respondents report strong perceived productivity gains

with the introduction of the AI-based tools. The logs show that *AI users* are writing more code than *AI non-users*.

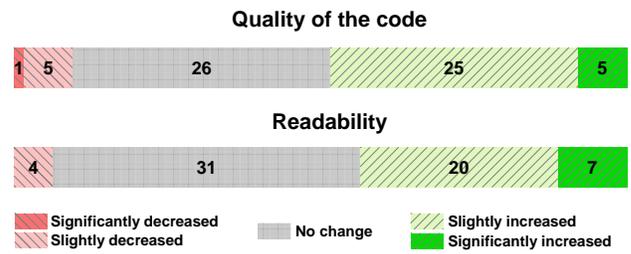
4.2 RQ2: Code Quality

The results for **code quality** are shown in Figure 2.

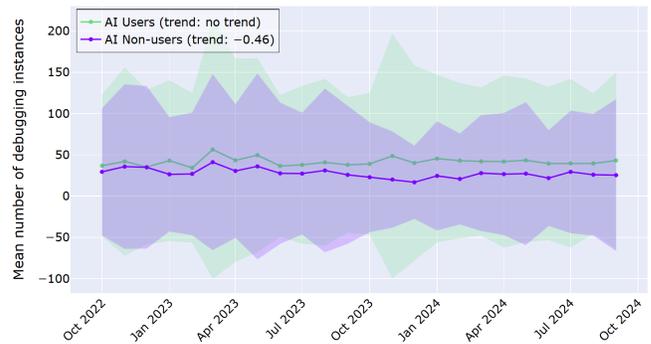
Survey and interviews (Figure 2a). When asked if the quality of their code increased because of using AI coding tools, 48.4% of respondents say that it slightly or significantly increased, while 9.7% say that it slightly or significantly decreased. For the readability of the code, the respective numbers are 43.5% and 6.5%, though 50% indicate that they did not observe a change. Overall, it can be seen that the view of the change in code quality is somewhat positive. However, P7 highlighted an interesting point about the relationship towards the quality of generated code:

“If something was done by AI, I’m actually more paranoid about it. I keep wondering, “Is this really correct?” I triple-check it, and even then, I still feel a bit uneasy.”

Logs (Figure 2b). We can see that for *AI users* the number of debugging events in logs has no statistically significant trend, meaning no noticeable change in the studied two years. In contrast, *AI*



(a) Responses to the survey questions whether *Quality of the code* and *Readability* decreased or increased due to using AI tools for coding.



(b) The mean *Number of debugging instances* per device in the IDE logs, for different months. *AI Users* show no trend, *AI Non-users* — a statistically significant decrease of -0.46 instances a month. Shaded regions represent ± 1 standard deviation from the monthly mean.

Figure 2: Results for RQ2: Code quality.

non-users exhibit a statistically significant decline in debugging activity, with an average decrease of approximately 0.46 debugging events per month ($p < 0.001$). This suggests that while debugging behavior remains stable among *AI users*, *AI non-users* are gradually engaging in less debugging over time.

Summary of RQ2. Survey respondents report moderately positive perceptions of code quality after using AI-based coding tools, though half perceive no change. Telemetry shows stable debugging activity among *AI users* and declining debugging activity among *AI non-users*

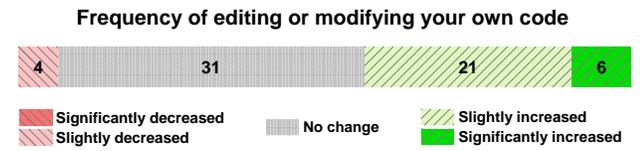
4.3 RQ3: Code Editing

The results for **code editing** are shown in Figure 3.

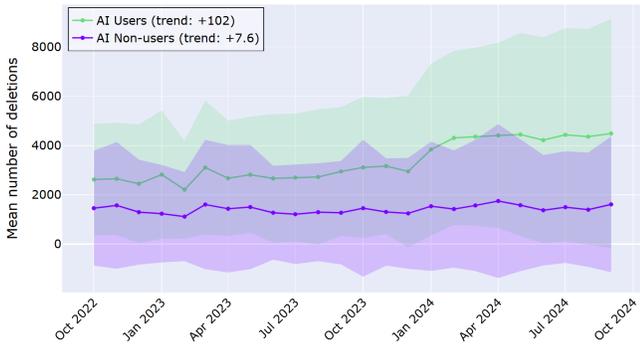
Survey and interviews (Figure 3a). Our results show that 43.5% of participants feel that the frequency of editing their code slightly or significantly increased with the usage of AI tools for coding, 6.5% feel like it decreased, and 50% believe it did not change.

P62, a system architect with more than 16 years of experience, noted that the regular use of AI tooling supports their code editing and reviewing practices:

“AI is like a second pair of eyes, offering pair programming benefits without social pressure — especially helpful for neurodivergent people. It’s not always watching,



(a) Responses to the survey question whether the *Frequency of editing or modifying your own code* decreased or increased due to using AI tools for coding.



(b) The mean *Number of deletions* per device in the IDE logs, for different months. *AI Users* show a statistically significant increasing trend of +102 deletions a month, *AI Non-users* — +7.6 deletions. Shaded regions represent ± 1 standard deviation from the monthly mean.

Figure 3: Results for RQ3: Code editing.

but I can call on it for code review and feedback when needed.”

Logs (Figure 3b). The analysis of deletions in the usage logs shows that, similar to typing, the number of deletions increased for both *AI users* and *AI non-users*, but it increased significantly faster for *AI users*, with an average rate of 102 deletions per month (whereas *AI non-users* showed an increase of about 7.6 deletions per month).

Summary of RQ3. Survey responses indicate that developers perceive little to no change in their code editing habits with AI, though some report a slight increase. Telemetry data reveals a clear rise in code deletions over time among AI users.

4.4 RQ4: Code Reuse

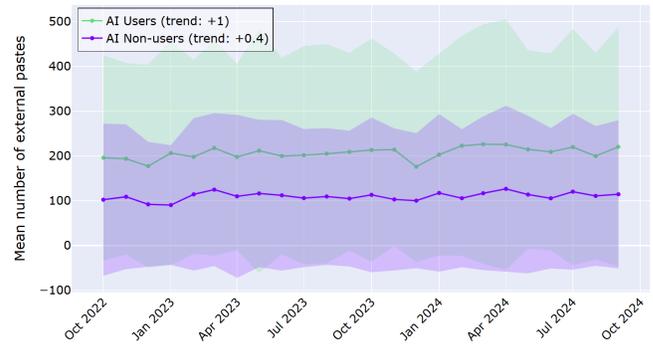
The results for **code reuse** are presented in Figure 4.

Survey and interviews (Figure 4a). 37.1% of respondents reported that their use of code from external sources slightly or significantly increased with the adoption of AI tools, while 19.4% said it decreased, and 43.5% observed no change. This distribution is notable, as code reuse might be expected to follow a more consistent trend. Instead, developers report divergent experiences. Representing the “decreasing” perspective, P39, a DevOps engineer with over 16 years of experience who discontinued AI tool usage, explained:

“For me, it’s better to take responsibility for what I did myself rather than adopt a third-party solution, try to understand it, and convince myself it’s correct.”



(a) Responses to the survey question whether the *Frequency of use of code from external sources* decreased or increased due to using AI tools for coding.



(b) The mean *Number of external pastes* per device in the IDE logs, for different months. *AI Users* show a statistically significant increasing trend of +1 paste a month, *AI Non-users* — +0.4 pastes. Shaded regions represent ± 1 standard deviation from the monthly mean.

Figure 4: Results for RQ4: Code reuse.

Logs (Figure 4b). In terms of actions of pasting from other sources in logs, we found a small, but statistically significant increase over time for both *AI users* and *AI non-users*. However, it is faster for *AI users* ($p = 0.03$). *AI users* show on average an increase of 1 paste per month, while *AI non-users* — 0.4 pastes per month.

Summary of RQ4. For code reuse, participant opinions on changes in frequency were relatively balanced, with no clear consensus. Telemetry reveals that paste actions from external sources increased over time for both groups, more prominently among AI users.

4.5 RQ5: Context Switching

Finally, the results for **context switching** are shown in Figure 5.

Survey and interviews (Figure 5a). Here, we can see that 25.8% of respondents indicated that their frequency of context switching slightly or significantly increased with the use of AI tools, 21% — that it decreased, and 53.2% reported no change. P62 expressed a thought that highlights how the introduction of AI decreased context switching in their workflow:

“I stopped switching contexts, saving a few seconds every time I would have googled something.”

Logs (Figure 5b). We can see that for *AI users*, there is a slight statistically significant increase in the number of context switches with an estimated growth of approximately 6.4 switches per month. For *AI non-users*, they had a slight statistically significant decrease of roughly 7.6 switches per month, as indicated by the negative

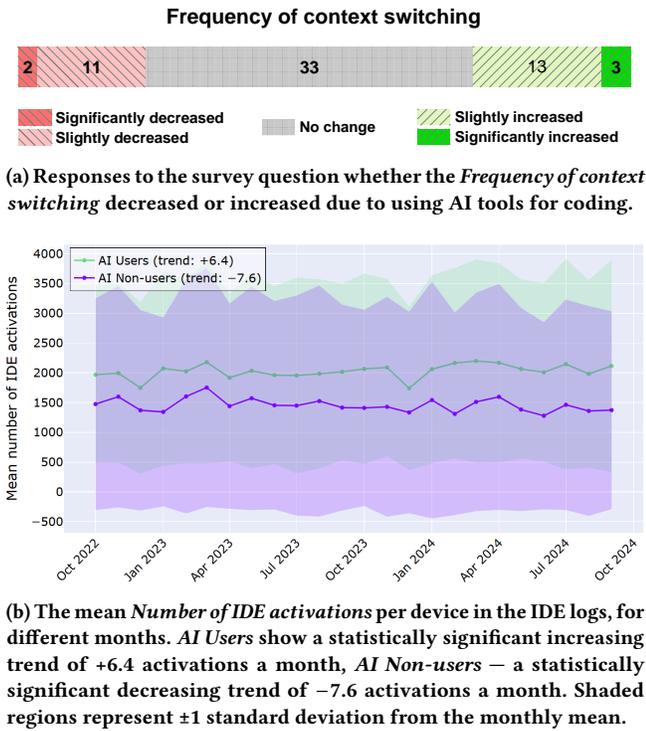


Figure 5: Results for RQ5: Context switching.

interaction effect. This bi-directional trend highlights the difference in AI's influence on developers' behavior.

Summary of RQ5. Context switching demonstrated approximately an equal number of participants reporting it increasing and decreasing with the usage of AI tools. As for the number of window switches in the log data, it showed a slight increase for AI users.

5 Discussion

To understand the big picture impact of AI on developer workflows, we examine self-reported perceptions and behavioral telemetry. Table 1 summarizes our findings across five key aspects.

5.1 Relation of Aspects with Each Other

Looking across the five workflow aspects, several inter-aspect patterns emerge that reveal how AI assistance influences not just isolated behaviors but broader development routines.

Perceived productivity gains among AI users are mirrored in telemetry through an increase in the amount of typing. Prior studies have similarly found that AI tools can accelerate code authoring [16, 27]. However, in logs, this increase in output is accompanied by a notable rise in code deletions. This suggests that productivity may come at a cost: developers produce more but also discard more, possibly due to iterative refinement of sometimes inefficient AI-generated suggestions [20]. Interestingly, no corresponding increase in debugging sessions was observed, which may indicate that errors are detected and addressed earlier in the workflow.

The discrepancy between increased deletions and stable debugging aligns with observations from the literature that developers treat AI suggestions as starting points rather than final answers [44].

Such patterns echo earlier findings that developers often spend substantial time verifying and revising AI outputs [23, 30]. Our survey shows that more than a third of respondents perceive the code quality and need for code editing as growing, but half report no change in these metrics, which is expected due to the tendency for over-reliance by some users of AI tooling [28, 31].

Code reuse and context switching also appear connected. In surveys, developers had no strong perception of change. Yet, in logs, we observe a slight but consistent increase in both pasting from external sources and activations reflecting window switches. While some literature anticipated a reduction in context switching due to in-IDE support [26], our data suggest that developers may still rely on external AI systems for validation or comparison. This hypothesis is supported by prior work documenting how developers use general-purpose tools like ChatGPT alongside IDE-integrated assistants [36].

5.2 Different Sources Tell Different Stories

Behavioral telemetry and perceptual survey responses address the research questions from complementary angles: telemetry shows what developers do, surveys show what developers recognize. Taken together, our findings reveal that AI tools reshape developer workflows in ways that are observable at the event level but not always apparent in holistic retrospection.

The telemetry data suggest that AI use is associated with increased activity volume in the IDE. Compared to *AI non-users*, *AI users* type and delete significantly more. In contrast, survey data portray a more reserved picture. Developers clearly report perceived productivity gains, but are more ambivalent across other aspects. For code quality, editing effort, reuse, and context switching, responses are often neutral, with roughly half of the participants reporting no observable change. These patterns echo prior work highlighting the disconnect between perceived and actual productivity [42, 46]. While these patterns are consistent with AI-mediated workflow changes, we acknowledge that early adopters may represent a distinct population whose behaviors do not fully generalize to all developers. Our findings characterize developers who consistently adopted AI over a sustained period.

AI-assisted workflows appear to reconfigure where effort occurs, rather than how much effort is perceived. This subtle reconfiguration may partly explain the field's mixed findings on AI impact [21, 37], and highlights the importance of complementing telemetry with perceptual data when evaluating developer-facing AI tools. Both perspectives are valid and complementary. Understanding AI's real-world impact requires examining both.

Overall, the combined picture points to a reframing of AI's role in development workflows: not as a direct substitute for effort or expertise, but as a modifier of rhythm, density, and focus of interaction. Rather than producing sweeping transformations visible to the user, AI tools appear to reshape the granularity and structure of everyday tasks. Recognizing this distinction is essential for both empirical evaluation and the design of future AI tools.

Table 1: An overview of the results of our study.

	Survey	IDE logs
Productivity	Respondents are highly positive, 82.3% say their productivity increased and 56.5% say that their coding time decreased.	The number of typed characters increases for both <i>AI users</i> and <i>AI non-users</i> , but grows much faster for <i>AI users</i> (+587 characters vs +75 for <i>AI non-users</i>).
Code quality	Respondents are relatively positive, 48.4% say that code quality increased and 43.5% see an increase in readability, with almost half not perceiving a change in both.	There is no significant change for <i>AI users</i> in terms of debugging sessions, while there is a slight negative trend of -0.46 sessions for <i>AI non-users</i> .
Code editing	43.5% of respondents feel that the frequency of code editing increased, but 50% see no change in it.	The number of deletions increases for both <i>AI users</i> and <i>AI non-users</i> , but grows faster for <i>AI users</i> (+102 deletions vs +7.6 for <i>AI non-users</i>).
Code reuse	Respondents are balanced in terms of using code from external sources, 37.1% believe that the frequency of use of code from external sources increased, 19.4% believe it decreased, and almost half see no change.	Pasting from external sources demonstrates slight positive trends for both groups, the trend for <i>AI users</i> (+1 paste per month) is higher than for <i>AI non-users</i> (+0.4).
Context switching	Respondents are split, 25.8% indicated that frequency of context switching increased, while 21% — that it decreased, and 53.2% reported no change.	There is a slight positive trend for IDE activations among <i>AI users</i> (+6.4), and a slight negative trend for <i>AI non-users</i> (-7.6).

5.3 Durability of Findings in Evolving AI Landscapes

Our study captures developer workflows during a critical transition period when AI assistance primarily operated through code completion and in-chat suggestions. The landscape has continued to evolve rapidly. Current agentic coding systems that autonomously generate larger code units are continuing to reshape development practice [6, 45]. This invites consideration of how our findings may evolve alongside advancing AI capabilities.

We argue that several core insights are durable. First, the gap between perception and behavior that we identify reflects fundamental human cognitive limitations in retrospective self-assessment [19], not properties of specific AI technologies. Whether code is generated via tab completion or autonomous agents, developers may similarly misjudge the extent to which their workflows have transformed. Second, the pattern of increased output coupled with increased deletion suggests a general dynamic of AI-assisted work: greater volume requires greater curation [23]. As AI generates more code, the editing and validation burden likely intensifies rather than disappears. Even as agents assume more authoring responsibility, developers must allocate their attention to specification, verification, and integration tasks.

Our work provides an empirical baseline for understanding how AI adoption reshaped workflows, enabling future research to measure whether subsequent AI paradigms amplify, reverse, or fundamentally alter these patterns. The complementary telemetry-survey methodology we demonstrate remains applicable for studying newer AI systems, even as the specific metrics must adapt to changing interaction modalities.

5.4 Design Implications for Future AI-Augmented IDEs

Our findings suggest that existing AI assistants are designed primarily for localized support (e.g., completing code) but lack awareness of broader workflow impacts. As development environments evolve toward more agentic systems and AI-native platforms, these limitations may become more pronounced or take different forms. To better support developers, future AI tools should not only generate code but also support context continuity, anticipate multi-step task flows, and provide visibility into where and how assistance is influencing developer behavior.

Furthermore, tools should support reflective practices—surfacing insights about the developer’s workflow evolution, rather than remaining silent partners. This becomes increasingly critical as AI assumes more autonomous roles: when agents generate substantial code, developers need explicit signals about changes in their verification effort, integration patterns, and overall workflow. For instance, telemetry-informed feedback on changes in reuse patterns or editing intensity could help developers regain situational awareness and avoid overreliance.

6 Threats to Validity

Construct validity. Grounded in prior work, we operationalized five workflow dimensions using in-IDE telemetry and survey responses. We deliberately included both behavioral and perceptual data, anticipating that they may represent different constructs. While telemetry ensures consistency and avoids recall bias, capturing fine-grained, moment-to-moment developer actions, it does not capture developer intent or task context. Conversely, survey data reflects holistic, retrospective perception but not actual behavior. To address these limitations, we complement data sources. In our case, these sources sometimes point in different directions. We do not treat this as a weakness, but as evidence that each source highlights

a different facet of AI-assisted workflows. To interpret our findings, readers should understand that behavioral telemetry and survey responses are not interchangeable measures of the same construct, but complementary windows into developer experience.

The study is based on telemetry from a single commercial IDE provider, JetBrains, and uses integration with a specific proprietary JetBrains AI assistant to define AI usage. Developers using other AI tools or AI tools outside the IDE may be misclassified as non-users, potentially underestimating true adoption. However, given the assistant's full integration and free trial availability, we assume that developers interested in using AI within the IDE would have at least tried the built-in assistant. Those who never activated it over two years, even once, are likely less inclined to adopt in-IDE AI support. Given that this was the only criterion for sampling the two separate groups, numerous statistically significant results indicate that the change is measurable, even if it does not fully generalize to users of all other AI tools.

Internal validity. Our data confirms that *AI users* are generally more active in the IDE than *AI non-users*, as evidenced by the fact that for most actions, *AI users* do more of them. It is likely that they differ not only in AI usage but also in baseline motivation, experience, or task profiles — all of which may influence AI adoption and observed behavior. However, our focus is on behavioral trends within each group over time, rather than cross-group causality. As such, we treat the distinction between *AI users* and *AI non-users* as descriptive rather than explanatory. For instance, when we observe that *AI users* show steeper increases in typed characters and deletions compared to *AI non-users*, this pattern could reflect AI's direct influence on coding intensity. Alternatively, it may indicate that early adopters maintain elevated activity levels regardless, and that AI integration occurred within an already active population. These interpretations cannot be fully disentangled without experimental assignment to conditions. The longitudinal nature of the analysis of trends supports the validity of comparisons within each group, while acknowledging that self-selection effects may shape between-group differences. Readers should interpret cross-group differences as associations reflecting both AI adoption and the unmeasured characteristics that predict adoption.

External validity. While the studied AI assistant offers general-purpose coding support comparable to other AI-based tools, we acknowledge that our findings may not generalize to all development environments or interface paradigms. At the same time, our data includes users of different IDEs, different programming languages, and different ecosystems. Moreover, the workflow dimensions we examine — productivity, code quality, code editing, code reuse, and context switching — are not tool-specific and remain relevant across a wide range of IDEs and development contexts. This supports the partial generalizability of our results beyond the specific assistant and platform studied.

7 Conclusion

We conducted a mixed-method study of AI adoption in IDEs, combining two-year longitudinal fine-grained telemetry from 800 professional developers, capturing over 151 million interaction events, with survey responses from 62 participants. Our findings reveal that

AI-assisted development workflows evolve in subtle but measurable ways over time. While *AI users* consistently report increased productivity and minimal changes in other dimensions, fine-grained telemetry traces show significant increases in writing and editing activity, as well as a rising trend in external code integration, and context switching. These results show that AI redistributes and reshapes development work, frequently in ways that elude developers' conscious awareness.

As AI becomes further integrated into the software development lifecycle, our results underscore the need for tooling and evaluation methods that account for both visible and more subtle, often overlooked shifts in workflow. Future research and tool design should focus not only on output metrics but also on workflow granularity, cognitive load, and user awareness to ensure that AI not only supports productivity but also aligns with and enhances core software engineering practices, rather than unintentionally reshaping them in opaque ways.

Data Availability

Supplementary materials, including the survey questionnaire and anonymized responses, interview script, and complete statistical analysis outputs, are publicly available in our supplementary materials repository [35].

Raw IDE telemetry logs cannot be released due to confidentiality agreements with our industry partner. All aggregated data and statistical outputs necessary to verify our findings are provided in the supplementary materials.

Researchers interested in accessing additional aggregated data or replicating our analysis methodology may contact the authors.

References

- [1] [n. d.]. ICC/ESOMAR codes and guidelines. <https://esomar.org/codes-and-guidelines>. Accessed: 2025-07-14.
- [2] Sven Amann, Sebastian Proksch, Sarah Nadi, and Mira Mezini. 2016. A study of Visual Studio usage in practice. In *2016 IEEE 23rd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*, Vol. 1. IEEE, 124–134.
- [3] Saulius Astromskis, Gabriele Bavota, Andrea Janes, Barbara Russo, and Massimiliano Di Penta. 2017. Patterns of developers behaviour: A 1000-hour industrial study. *Journal of Systems and Software* 132 (2017), 85–97.
- [4] Joel Becker, Nate Rush, Elizabeth Barnes, and David Rein. 2025. Measuring the impact of early-2025 AI on experienced open-source developer productivity. *arXiv preprint arXiv:2507.09089* (2025).
- [5] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374* (2021).
- [6] Cognition AI. 2024. Devin: AI Software Engineer. <https://www.cognition.ai/blog/introducing-devin>. Accessed: 2025-01-15.
- [7] Mariana Coutinho, Lorena Marques, Anderson Santos, Marcio Dahia, Cesar França, and Ronnie de Souza Santos. 2024. The role of generative AI in software development productivity: A pilot case study. In *Proceedings of the 1st ACM International Conference on AI-Powered Software*. 131–138.
- [8] Arghavan Moradi Dakhel, Vahid Majdinasab, Amin Nikanjam, Foutse Khomh, Michel C Desmarais, and Zhen Ming Jack Jiang. 2023. GitHub Copilot AI pair programmer: Asset or liability? *Journal of Systems and Software* 203 (2023), 111734.
- [9] Kostadin Damevski, Hui Chen, David Shepherd, and Lori Pollock. 2016. Interactive exploration of developer interaction traces using a hidden Markov model. In *Proceedings of the 13th International Conference on Mining Software Repositories*. 126–136.
- [10] Maria del Rio-Chanona, Nadzeya Laurentsyeve, and Johannes Wachs. 2023. Are large language models a threat to digital public goods? Evidence from activity on Stack Overflow. *arXiv preprint arXiv:2307.07367* (2023).
- [11] Fabrizio Dell'Acqua, Charles Ayoubi, Hila Lifshitz, Raffaella Sadun, Ethan Mollick, Lilach Mollick, Yi Han, Jeff Goldman, Hari Nair, Stewart Taub, et al. 2025. *The*

- cybernetic teammate: A field experiment on generative AI reshaping teamwork and expertise. Technical Report. National Bureau of Economic Research.
- [12] Prem Devanbu, Thomas Zimmermann, and Christian Bird. 2016. Belief & evidence in empirical software engineering. In *Proceedings of the 38th international conference on software engineering*, 108–119.
 - [13] Nicole Forsgren, Margaret-Anne Storey, Chandra Maddila, Thomas Zimmermann, Brian Houck, and Jenna Butler. 2021. The SPACE of developer productivity: There's more to it than you think. *Queue* 19, 1 (2021), 20–48.
 - [14] GitHub and OpenAI. [n. d.]. GitHub Copilot. <https://github.com/features/copilot/> Accessed: 2025-07-14.
 - [15] William Harding. 2025. *AI Copilot code quality: Evaluating 2024's increased defect rate via code quality metrics*. Technical Report. Alloy.dev Research. https://www.gitclear.com/ai_assistant_code_quality_2025_research Accessed: 2025-07-14.
 - [16] Saki Imai. 2022. Is GitHub Copilot a substitute for human pair-programming? An empirical study. In *Proceedings of the ACM/IEEE 44th International Conference on Software Engineering: Companion Proceedings*. 319–321.
 - [17] JetBrains. [n. d.]. JetBrains AI Assistant. <https://plugins.jetbrains.com/plugin/22282-jetbrains-ai-assistant> Accessed: 2025-07-14.
 - [18] JetBrains. 2024. The State of Developer Ecosystem 2024. <https://www.jetbrains.com/lp/devecosystem-2024/>. Accessed 2025-07-14.
 - [19] Daniel Kahneman. 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
 - [20] Sila Lertbanjongngam, Bodin Chinthanet, Takashi Ishio, Raula Gaikovina Kula, Pattara Leelaprute, Bundit Manaskasemsak, Arnon Rungsawang, and Kenichi Matsumoto. 2022. An empirical evaluation of competitive programming AI: A case study of Alphacode. In *2022 IEEE 16th International Workshop on Software Clones (IWSC)*. IEEE, 10–15.
 - [21] Jenny T Liang, Chenyang Yang, and Brad A Myers. 2024. A large-scale survey on the usability of AI programming assistants: Successes and challenges. In *Proceedings of the 46th IEEE/ACM international conference on software engineering*, 1–13.
 - [22] Roberto Minelli, Andrea Mocci, and Michele Lanza. 2015. I know what you did last summer—an investigation of how developers spend their time. In *2015 IEEE 23rd international conference on program comprehension*. IEEE, 25–35.
 - [23] Hussein Mozannar, Gagan Bansal, Adam Fourney, and Eric Horvitz. 2024. Reading between the lines: Modeling user behavior and costs in AI-assisted programming. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–16.
 - [24] Emerson Murphy-Hill, Ciera Jaspan, Caitlin Sadowski, David Shepherd, Michael Phillips, Collin Winter, Andrea Knight, Edward Smith, and Matthew Jorde. 2019. What predicts software developers' productivity? *IEEE Transactions on Software Engineering* 47, 3 (2019), 582–594.
 - [25] OpenAI. 2022. ChatGPT: Mar 14 version. <https://openai.com/blog/chatgpt> Accessed: 2025-07-14.
 - [26] Ruchika Pandey, Prabhat Singh, Raymond Wei, and Shailla Shankar. 2024. Transforming software development: Evaluating the efficiency and challenges of GitHub Copilot in real-world projects. *arXiv preprint arXiv:2406.17910* (2024).
 - [27] Sida Peng, Eirini Kalliamvakou, Peter Cihon, and Mert Demirel. 2023. The impact of AI on developer productivity: Evidence from GitHub Copilot. *arXiv preprint arXiv:2302.06590* (2023).
 - [28] James Prather, Brent N Reeves, Paul Denny, Brett A Becker, Juho Leinonen, Andrew Luxton-Reilly, Garrett Powell, James Finnie-Ansley, and Eddie Antonio Santos. 2023. "It's weird that it knows what I want": Usability and interactions with Copilot for novice programmers. *ACM Transactions on Computer-Human Interaction* 31, 1 (2023), 1–31.
 - [29] Steven I Ross, Fernando Martinez, Stephanie Houde, Michael Muller, and Justin D Weisz. 2023. The programmer's assistant: Conversational interaction with a large language model for software development. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*. 491–514.
 - [30] Priyam Sahoo, Saurabh Pujar, Ganesh Nalawade, Richard Genhardt, Louis Mandel, and Luca Buratti. 2024. Ansible Lightspeed: A code generation service for IT automation. In *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering*. 2148–2158.
 - [31] James Schaffer, John O'Donovan, James Michaelis, Adrienne Raglin, and Tobias Höllerer. 2019. I can do better than your AI: expertise and explanations. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*. 240–251.
 - [32] SciPy. 2024. *scipy.stats.bartlett — SciPy v1.13.0 Manual*. <https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.bartlett.html> Accessed: 2025-07-14.
 - [33] SciPy. 2024. *scipy.stats.kstwo — SciPy v1.13.0 Manual*. <https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.kstwo.html> Accessed: 2025-07-14.
 - [34] Agnia Sergeyyuk, Yaroslav Golubev, Timofey Bryksin, and Iftekhar Ahmed. 2025. Using AI-based coding assistants in practice: State of affairs, perceptions, and ways forward. *Information and Software Technology* 178 (2025), 107610.
 - [35] Agnia Sergeyyuk, Eric Huang, Dariia Karaeva, Anastasiia Serova, Yaroslav Golubev, and Iftekhar Ahmed. 2025. Supplementary materials. <https://doi.org/10.5281/zenodo.18220544>. Accessed: 2026-01-12.
 - [36] Agnia Sergeyyuk, Ekaterina Koshchenko, Ilya Zakharov, Timofey Bryksin, and Maliheh Izadi. 2024. The Design Space of in-IDE Human-AI Experience. *arXiv preprint arXiv:2410.08676* (2024).
 - [37] Fangchen Song, Ashish Agarwal, and Wen Wen. 2024. The impact of generative AI on collaborative open-source software development: Evidence from GitHub Copilot. *arXiv preprint arXiv:2410.02091* (2024).
 - [38] Stack Overflow. 2024. Stack Overflow developer survey 2024. <https://survey.stackoverflow.co/2024/ai>. Accessed 2025-07-14.
 - [39] statsmodels. 2024. *Mixed Linear Model — statsmodels v0.14.0 Manual*. https://www.statsmodels.org/stable/mixed_linear.html Accessed: 2025-07-14.
 - [40] Benyamin Tabarsi, Heidi Reichert, Ally Limke, Sandeep Kuttal, and Tiffany Barnes. 2025. LLMs' reshaping of people, processes, products, and society in software development: A comprehensive exploration with early adopters. *arXiv preprint arXiv:2503.05012* (2025).
 - [41] Thiago S Vaillant, Felipe Deveza de Almeida, Paulo Anselmo Neto, Cuiyun Gao, Jan Bosch, and Eduardo Santana de Almeida. 2024. Developers' perceptions on the impact of ChatGPT in software development: A survey. *arXiv preprint arXiv:2405.12195* (2024).
 - [42] Priyan Vaithilingam, Tianyi Zhang, and Elena L Glassman. 2022. Expectation vs. experience: Evaluating the usability of code generation tools powered by large language models. In *Chi conference on human factors in computing systems extended abstracts*. 1–7.
 - [43] Thomas Weber, Maximilian Brandmaier, Albrecht Schmidt, and Sven Mayer. 2024. Significant productivity gains through programming with large language models. *Proceedings of the ACM on Human-Computer Interaction* 8, EICS (2024), 1–29.
 - [44] Justin D Weisz, Michael Muller, Steven I Ross, Fernando Martinez, Stephanie Houde, Mayank Agarwal, Kartik Talamadupula, and John T Richards. 2022. Better together? An evaluation of AI-supported code translation. In *Proceedings of the 27th International Conference on Intelligent User Interfaces*. 369–391.
 - [45] John Yang, Carlos E Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. 2024. SWE-agent: Agent-computer interfaces enable automated software engineering. *Advances in Neural Information Processing Systems* 37 (2024), 50528–50652.
 - [46] Albert Ziegler, Eirini Kalliamvakou, X Alice Li, Andrew Rice, Devon Rifkin, Shawn Simister, Ganesh Sittampalam, and Edward Aftandilian. 2022. Productivity assessment of neural code completion. In *Proceedings of the 6th ACM SIGPLAN International Symposium on Machine Programming*. 21–29.